

Exhibit 1

Redacted in its Entirety

Exhibit 2

Redacted in its Entirety

Exhibit 3

Redacted in its Entirety

Exhibit 4

Redacted in its Entirety

Exhibit 5

Redacted in its Entirety

Exhibit 6

Report

ERCOT Market Mechanisms

**Prepared by
Shams Siddiqi, Ph.D.**

May 6, 2022

Table of Contents

Table of Contents	i
1. Executive Summary and Outline of Report	1
1.1 Executive Summary	1
1.2 List of Documents Considered	1
1.3 Compensation	2
1.4 Outline of Report	2
2. Overview of the ERCOT Market	4
2.1 Overview of ERCOT Market	4
2.2 Unique Features of ERCOT Nodal Market	5
2.3 Market Participants and Their Roles in ERCOT Market	7
3. ERCOT Market Mechanics	11
3.1 Security Constrained Economic Dispatch (SCED) and Locational Marginal Price (LMP)	11
3.2 Settlement Point Price (SPP)	13
3.3 ERCOT Ancillary Service Market	13
3.4 ERCOT Day-Ahead Market (DAM)	14
4. Bilateral Market Mechanics	20
4.1 Bilateral Contracting in ERCOT	20
4.2 Settlement of Bilateral Contracts	20
5. Generation and Load Netting Provision	22
6. Load Participation in ERCOT Markets	23
6.1 Load Participation in Energy Market – Price Response	24
6.2 Load Participation in Energy Market – Energy Arbitrage	24
6.3 Load Participation in Emergency Response Service (ERS)	25
6.4 Load Participation in ERCOT Ancillary Service Market	26
6.5 Load Participation as Controllable Load Resource (CLR)	27
6.6 Load Participation as Non-Controllable Load Resource (NCLR)	28
7. Critique of Mr. McCamant’s Report	29
8. Conclusion	31
8.1 ERCOT Market has long history of Load Participation	31

8.2	Mr. McCamant's Report has Erroneous Assertion	31
8.3	Right to Supplement or Amend based on Discovery.....	31
9.	Resume of Shams Siddiqi, Ph.D.....	32
10.	Industry Terms and Acronyms	37

However, if the 100 MW CLR submit a RTM Energy Bid of \$40/MWh into SCED, then that CLR bid would not be awarded and the CLR would be dispatched by SCED to consumer 0 MW, A1 dispatched to 550 MW, A2 dispatched to 50 MW, B1 dispatched to 800 MW, and B2 dispatched to 0 MW. The resulting LMPs are \$20/MWh at Bus A and \$50/MWh at Bus B.

Nodal, Zonal, and Hub LMPs

In the example above, Bus A and Bus B are what are called Nodes in the system. The prices calculated above for Bus A and Bus B are called Nodal LMPs since the prices are tied to specific Nodes. However, loads are settled on a Zonal basis. To illustrate this, assume that the system in Figure 4 has only one Load Zone consisting of load at Bus A and Bus B.

The Load Zone LMP is calculated as the load-weighted average of all bus LMPs within the Load Zone. Thus, for the SCED Nodal LMPs of \$20/MWh at Bus A and \$70/MWh at Bus B, the Load Zone LMP = $(20 \times 500 + 70 \times 1000) / 1500 = \$53.33/\text{MWh}$. Even though A1 and LoadA are both located at Bus A, there is congestion exposure between them due to Load Zone price averaging – i.e. for settlement purposes, LoadA is at the virtual Load Zone location and not at its physical location of Bus A.

The Hub LMP is calculated the simple average of all bus LMPs that constitute the Hub. If an ISO Hub is defined as to include all buses in the ISO system, then in the example above, the ISO Hub LMP = $(20 + 70) / 2 = \$45/\text{MWh}$.

3.2 Settlement Point Price (SPP)

Settlement Point Price is defined as “A price calculated for a Settlement Point for each Settlement Interval using LMP data and the formulas detailed in Section 4.6, DAM Settlement, and Section 6.6, Settlement Calculations for the Real-Time Energy Operations.”

Real-Time energy Settlements use Real-Time Settlement Point Prices that are calculated for Resource Nodes, Load Zones, and Hubs. For each Security-Constrained Economic Dispatch (SCED) Locational Marginal Price (LMP) calculated at each Settlement Point in the SCED process, an administrative price floor of -\$251/MWh will be applied to Real-Time Settlement Point Prices after adding the sum of the Real-Time On-Line Reliability Deployment Price Adders and the Real-Time On-Line Reserve Price Adder.

3.3 ERCOT Ancillary Service Market

ERCOT operates a Day-Ahead Market (clearing price auction) and is provider of last resort for:

- Regulation Up and Down Services (Reg-Up or RUS and Reg-Down or RDS)
- Responsive Reserves Service (RRS)
- Non-Spinning Reserves Service (NSRS)

Ancillary Service capacity obligations are allocated to QSEs on a Load Ratio Share basis. Entities may self-arrange Reg-Up, Reg-Down, RRS, and NSRS.

Regulation Service provides capacity that can respond to signals from ERCOT within three to five seconds to respond to changes from scheduled system frequency. A QSE control system must be capable of receiving digital control signals from ERCOT's control system, and of directing its Resources to respond to the control signals, in an upward and downward direction. Resources must be capable of delivering the full amount of regulating capacity offered to ERCOT within five minutes.

Responsive Reserve Service (RRS) provides capacity that can arrest frequency decay after a frequency event (e.g. loss of large generators) and may be provided by: unloaded Generation Resources that are On-Line, Load Resources controlled by high-set under-frequency relays or controllable like Generation Resources, and Fast Frequency Response (FFR – introduced on March 1, 2020) that responds in 15 cycles when frequency drops below 59.85Hz. Generation Resource must be capable of ramping to RRS capacity within ten minutes and can provide up to 20% of thermal unit High Sustainable Limit (HSL) and must be less than or equal to 10 times Resource's Emergency Ramp Rate. ERCOT procures enough RRS to arrest frequency prior to reaching 59.4Hz with the loss of 2,750MW.

Non-Spinning Reserve Service (NSRS) provides capacity that can replenish deployed RRS and be used for net load (load minus intermittent resource output) forecast errors and is provided by using Generation Resources, whether On-Line or Off-Line, capable of being synchronized and ramped to a specified output level within 30 minutes, running at a specified output level for at least one hour and Load Resources capable of being interrupted within 30 minutes and remaining interrupted for at least one hour.

3.4 ERCOT Day-Ahead Market (DAM)

ERCOT Day-Ahead Market (DAM) is a voluntary financial market that provides Market Participant a chance to trade day-ahead for the following Operating Day – i.e., awards of DAM PTP Obligations, energy bids/offers and Three-Part Supply Offers (TPSO) from Generation Resources are not physically binding but do create financial obligations. However, Ancillary Service (AS) awards are physically binding. DAM co-optimizes AS, Energy, and PTP Obligations. DAM allows energy trading, purchase of PTP Obligations that are settled on Real-time prices, energy bids/offers, AS offers, and Three-Part Energy offers by resources (start-up cost and minimum-energy cost parts are capped at Resource Category Startup and Minimum-Energy Generic Caps or Resource-specific Verifiable Costs).

Ancillary Service Offer

By 10am³ in Day-Ahead, QSE may submit Resource-specific AS Offers for DAM and generators may offer the same capacity for any or all of AS products simultaneously with any Energy Offer Curves (EOC) from that Resource in DAM. QSE may also submit AS Offers in Supplemental Ancillary Service Market (SASM). Offers of more than one AS product from one Generation Resource may be inclusive or exclusive of each other and of any EOC. By 10am in Day-Ahead, QSE may submit Load Resource-specific AS Offers for DAM and may offer the same capacity for any or all of those AS products that it is qualified to provide simultaneously. Offers of more than one AS

³ All times specified in this Report are Central Prevailing Time.

product from one Load Resource may be inclusive or exclusive of each other.

Three-Part Supply Offer (TPSO)

Three-Part Supply Offer consists of Startup Offer (SUO), Minimum-Energy Offer (MEO) & Energy Offer Curve (EOC). DAM uses all three parts of TPSO and also uses EOCs submitted without SUO and without MEO. RUC only uses SUO and MEO for RUC commitments, but EOC may be used in Settlement to claw back some or all of a RUC-committed Resource's energy payments. EOC may also be used by SCED in Real-Time Operations. QSE that submits EOC without also submitting SUO and MEO is considered not to be offering the Resource into the RUC, but that does not prevent the Resource from being committed in the RUC process like any other Resource that does not submit an offer in the RUC.

Energy Offer Curve (EOC)

EOC represents the QSE's willingness to sell energy at or above a certain price and at a certain quantity in DAM or its willingness to be dispatched by SCED. QSE may submit Resource-specific EOC which are bounded in DAM for each Operating Hour by LSL and HSL of the Resource specified in the COP and bounded in SCED by LSL and HSL of the Resource as shown by telemetry. EOCs remain active for the offered period until either selected by ERCOT or automatically inactivated at offer expiration time selected by QSE. For any Operating Hour, QSE may submit or change EOCs in the Adjustment Period and may withdraw an EOC if Output Schedule is submitted for intervals when EOC is withdrawn or Resource is forced Off-Line and changes Resource Status appropriately.

For any Operating Hour that is a RUC-Committed Interval or a DAM-Committed Interval for a Resource, a QSE for that Resource may not change a SUO or MEO. If a valid EOC or an Output Schedule does not exist for a Resource that has a status of On-Line at the end of the Adjustment Period, then ERCOT shall notify the QSE and set the Output Schedule equal to the then current telemetered output of the Resource until an Output Schedule or Energy Offer Curve is submitted in a subsequent Adjustment Period.

DAM Energy-Only Offer

QSE must submit any DAM Energy-Only Offer Curves by 10am for the effective DAM. DAM Energy-Only Offer Curve represents the QSE's willingness to sell energy at or above a certain price and at a certain quantity at a specific Settlement Point in DAM. DAM Energy-Only Offer Curves are not Resource-specific.

DAM Energy Bid

QSE must submit any DAM Energy Bids by 10am. DAM Energy Bid represents the QSE's willingness to buy energy at or below a certain price and at a certain quantity at a specific Settlement Point in the DAM and must include:

- The buying QSE;
- The Settlement Point;
- Fixed quantity block, variable quantity block, or curve indicator;
- If a fixed quantity block, the single price (in \$/MWh) and single quantity (in MW) for all hours bid in that block, which may clear at a Settlement Point Price greater than the bid price for that block;

- If a variable quantity block, the single price (in \$/MWh) and single “up to” quantity (in MW) contingent on the purchase of all hours bid in that block; and
- If a curve, a monotonically decreasing energy bid curve for both price (in \$/MWh) and quantity (in MW) with no more than 10 price/quantity pairs.
- The first and last hour of the bid; and
- The expiration time and date of the bid.
- Minimum amount for each DAM Energy Bid is one MW.

DAM Clearing Process

DAM uses a multi-hour mixed integer programming algorithm to maximize bid-based revenues minus the offer-based costs over the Operating Day, subject to security and other constraints, and ERCOT Ancillary Service procurement requirements. The bid-based revenues include revenues from DAM Energy Bids and Point-to-Point (PTP) Obligation bids. The offer-based costs include costs from SUO, MEO, and EOC of any Resource that submitted TPSO, DAM Energy-Only Offers and AS Offers.

Security constraints specified to prevent DAM solutions that would overload the elements of ERCOT Transmission Grid include:

- Transmission constraints – transfer limits on energy flows through the ERCOT Transmission Grid, e.g., thermal or stability limits. These limits must be satisfied by the intact network and for certain specified contingencies. These constraints may represent:
- Thermal constraints – protect Transmission Facilities against thermal overload.
- Generic constraints – protect the ERCOT Transmission Grid against transient instability, dynamic stability or voltage collapse.
- Power flow constraints – the energy balance at required Electrical Buses in the ERCOT Transmission Grid must be maintained.
- Resource constraints – the physical and security limits on Resources that submit Three-Part Supply Offers:
- Resource output constraints – the Low Sustained Limit (LSL) and High Sustained Limit (HSL) of each Resource; and
- Resource operational constraints – includes minimum run time, minimum down time, and configuration constraints.
- Other constraints –
 - Linked offers – the DAM may not select any one part of that Resource capacity to provide more than one Ancillary Service or to provide both energy and an Ancillary Service in the same Operating Hour. The DAM may, however, select part of that Resource capacity to provide one Ancillary Service and another part of that capacity to provide a different Ancillary Service or energy in the same Operating Hour, provided that a Generation Resource may not offer, and the DAM may not select, linked Energy and Off-Line Non-Spinning Reserve (Non-Spin) Ancillary Service Offers in the same Operating Hour.
 - The sum of the awarded Ancillary Service capacities for each Resource must be within the Resource limits specified in the Current Operating Plan (COP) and Section 3.18, Resource Limits in Providing Ancillary Service, and the Resource Parameters as described in Section 3.7, Resource Parameters.
 - Block Ancillary Service Offers for a Load Resource – blocks will not be cleared unless the entire quantity block can be awarded. Because block Ancillary Service Offers cannot set

the Market Clearing Price for Capacity (MCPC), a block Ancillary Service Offer may clear below the Ancillary Service Offer price for that block.

- Block DAM Energy Bids, DAM Energy-Only Offers, and PTP Obligation bids – blocks will not be cleared unless the entire time and/or quantity block can be awarded. Because quantity block bids and offers cannot set the Settlement Point Price, a quantity block bid or offer may clear in a manner inconsistent with the bid or offer price for that block.
- Combined Cycle Generation Resources – The DAM may commit a Combined Cycle Generation Resource in a time period that includes the last hour of the Operating Day only if that Combined Cycle Generation Resource can transition to a shutdown condition in the DAM Operating Day.

Ancillary Service needs for each Ancillary Service include the needs specified in the Ancillary Service Plan that are not part of the Self-Arranged Ancillary Service Quantity and that must be met from available DAM Ancillary Service Offers while co-optimizing with DAM Energy Offers. ERCOT may not buy more of one Ancillary Service in place of the quantity of a different service. ERCOT determines the appropriate Load distributions to allocate offers, bids, and source and sink of CRRs at a Load Zone across the Electrical Buses that are modeled with Load in that Load Zone. ERCOT allocates offers, bids, and source and sink of CRRs at a Hub using the distribution factors specified in the definition of that Hub. Resource that has a Three-Part Supply Offer cleared in the DAM may be eligible for Make-Whole Payment of the Startup Offer and Minimum Energy Offer submitted by the QSE representing the Resource. DAM energy offer awards create financial short (load) liabilities in RTM and DAM energy bid awards create financial long (supply) credits in RTM for awarded amounts.

Communicating DAM Results

As soon as practicable, but no later than 1:30pm in the Day-Ahead, ERCOT shall notify the parties to each cleared DAM transaction (e.g., the buyer and the seller) of the results of the DAM as follows:

- a. Awarded Ancillary Service Offers, specifying Resource, MW, Ancillary Service type, and price, for each hour of the awarded offer;
- b. Awarded energy offers from Three-Part Supply Offers and from DAM Energy-Only Offers, specifying Resource (except for DAM Energy-Only Offers), MWh, Settlement Point, and Settlement Point Price, for each hour of the awarded offer;
- c. Awarded DAM Energy Bids, specifying MWh, Settlement Point, and Settlement Point Price for each hour of the awarded bid; and
- d. Awarded PTP Obligation Bids, number of PTP Obligations in MW, source and sink Settlement Points, and price for each Settlement Interval of the awarded bid.

Example of Day-Ahead Market Mechanics

To describe DAM, consider again the simple two-bus (represented by the bold vertical lines) system with Bus A having two generators (represented by circles) A1 and A2 and Bus B having two generators B1 and B2 and two identical transmission lines T1 and T2 connecting Bus A and Bus B as shown in the figure below and demand is bid at the Load Zone as shown.




Figure 5. Simple Two-Bus System DAM Example

In addition to bids and offers above, there is a 100 MW PTP Obligation bid from A to B at \$65/MW/hour. Since loads bid at Load Zones, how are these bids modeled? ERCOT specifies Load Distribution Factors (LDF) to distribute Load Zone energy bids to various buses in DAM based on expected real-time distribution of metered load – assume LDFs of 1/3 at A and 2/3 at B. The Load Zone LMP for a Load Zone with n load buses can be calculated as:

$$LMP_{LZ} = \sum_{k=1,n} LDF_k * LMP_k$$

How does the ISO clear this market? What are LMPs and settlement statements?

This is more complex than the SCED problem since there is no physical load that ERCOT is trying to serve but rather energy bids and PTP Obligation bids. Since the PTP Obligation bid is greater than possible LMP differences between A and B, the PTP Obligation bid is fully awarded for 100 MW. That leaves no remaining transfer capability from A to B. Thus, the DAM awards are as follows:

A1 is awarded 500 MW, A2 is awarded 0, B1 is awarded 850 MW, B2 is awarded 150 MW, PTP Obligation from A to B is awarded 100 MW, and all load bids are awarded 1,500 MW.

The DAM clearing prices are: \$10/MWh at A, \$70/MWh at B, \$50/MWh at the Load Zone, \$60/MW/hour for PTP Obligation from A to B.

The DAM settlement, assuming that a CRR Owner bought 100 MW CRR from A to B in a CRR Auction, is given by (negative dollar amounts are credits to Market Participants and positive dollar amounts are payments from Market Participants):

WGR	Wind-powered Generation Resource
WGRPP	Wind-powered Generation Resource Production Potential
WRUC	Weekly Reliability Unit Commitment
WSL	Wholesale Storage Load



Shams Siddiqi

Exhibit 7

Redacted in its Entirety

Exhibit 8

Redacted in its Entirety


Exhibit 9

Redacted in its Entirety

Exhibit 10

Redacted in its Entirety

Exhibit 11



BearBox

@BearBox_io

Rugged, high-quality Bitcoin mining mobile infrastructure. Send us a DM or email contact@bearbox.io for more info!

New Orleans, LA Joined September 2018

75 Following 391 Followers

Tweets

Tweets & replies

Media

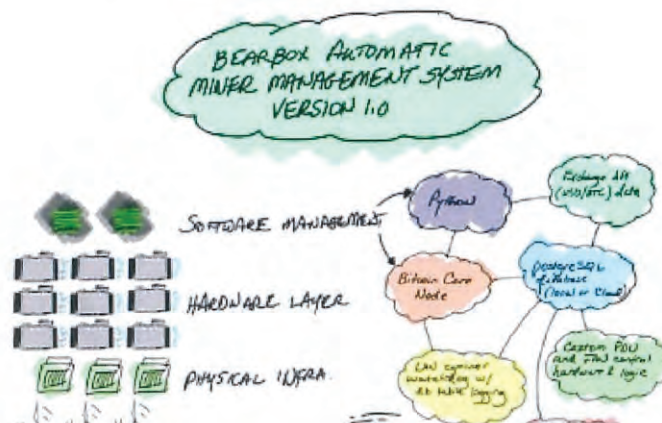
Likes

You Retweeted



Austin Storms @austorms · May 22, 2020

Casually pitched the idea of utilizing Bitcoin mining as LaaS (load-as-a-service) to a PM @Entergy early last year.



EXHIBIT

tabbles

D 68

 You Retweeted



Austin Storms @austorms · May 16, 2019

FYI - If you're looking for a great little temperature (or other) sensor platform for your mining operation, the AVTECH Room Alert 4E and others are really fantastic and output readings to JSON locally.

1

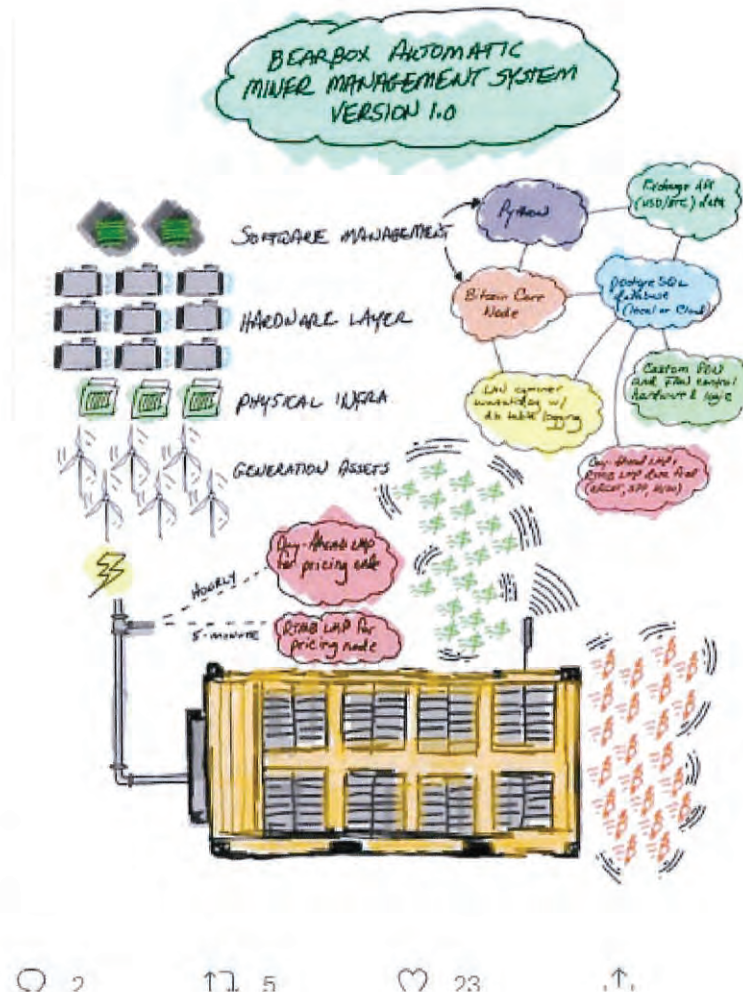
↕ 1

5

[Show this thread](#)

BearBox @BearBox_io · Jun 24, 2019

Converting generated ⚡⚡ into a liquid asset like [SBTC](#) (instead of storing it locally as an illiquid asset, re: lithium ion battery storage solutions) is a absolute no-brainer.



Q 2

↑ 17. 5

♡ 23

111

← **BearBox** 
58 Tweets

↻ You Retweeted



Austin Storms @austorms · Jun 16, 2019

The @whatsminer M20 is the best Bitcoin ASIC on the market right now, IMHO:

- \$1,699 (\$37.55/TH)
- 45TH/s (8nm)
- 2,160w (48w/TH - still within IEC 60320, C13/C14 @ 240/415v L-N)
- Form factor (same as last gen, diff. for high-density deployments and physical rack re-config.)

💬 3

↻ 5

❤️ 39



[Show this thread](#)

↻ You Retweeted



ODELL @ODELL · May 11, 2019

If you started buying \$25 worth of bitcoin every week at the top of the market, December 15th 2017, you would be in profit as of this week.

Sats stacked: 31,050,000

Cost: \$1850

Current value: \$1980

stackingsats.com [#stackingsats](#)



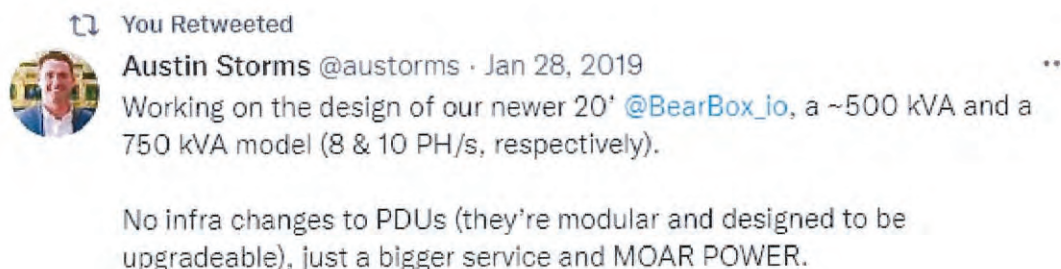
tesla.com

Electric Cars, Solar & Clean Energy | Tesla

Tesla is accelerating the world's transition to sustainable energy with



[Show this thread](#)



← **BearBox** 
58 Tweets

↻ You Retweeted



Austin Storms @austorms · Jan 28, 2019

Working on the design of our newer 20' @BearBox_io, a ~500 kVA and a 750 kVA model (8 & 10 PH/s, respectively).


No infra changes to PDUs (they're modular and designed to be upgradeable), just a bigger service and MOAR POWER.



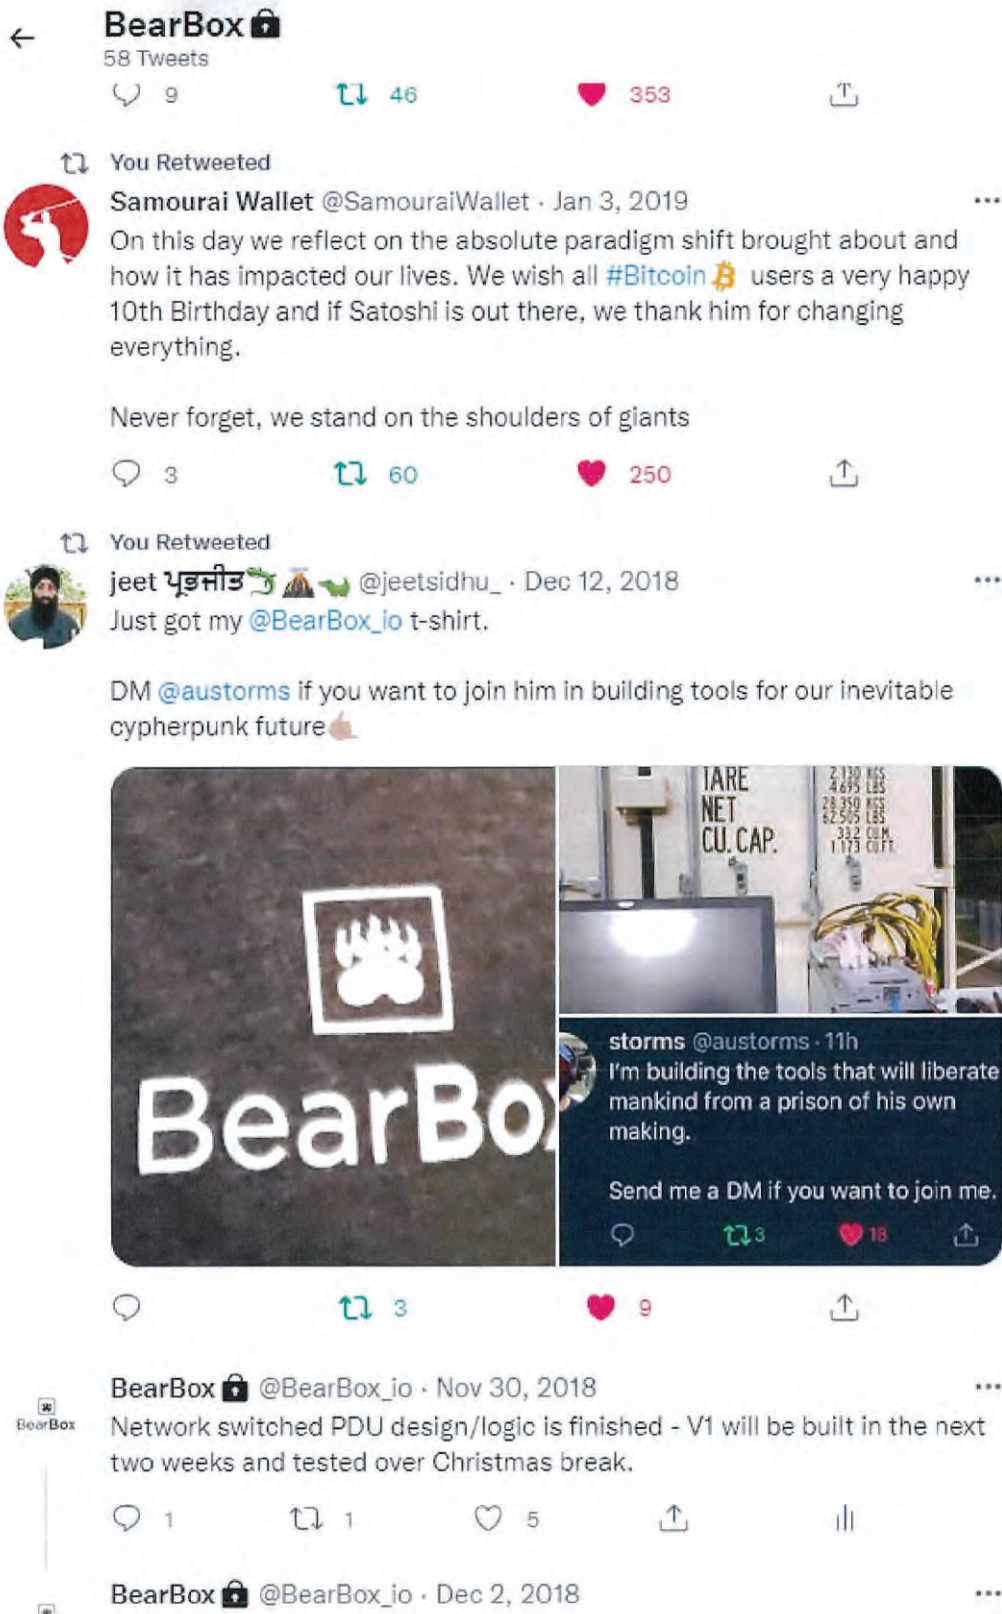
↻ You Retweeted





Austin Storms @austorms · Dec 27, 2018

It's simple, really. #Bitcoin 







 **BearBox**  @BearBox_io · Dec 2, 2018 ...
Update: created a *very* basic GUI for PDU remote relay control yesterday.



1 2

 **BearBox**  @BearBox_io · Dec 2, 2018 ...
*because this scales and existing applications don't

4

 You Retweeted
 **Power Mining** @PowerMiningFarm · Dec 2, 2018 ...
1/2 Seems that soon Bitcoin network will experience the biggest difficulty drop within the last year of almost 15%.
That clearly tells that the past 2 weeks a lot of Bitcoin miners has switched



BearBox @BearBox_io · Nov 23, 2018

...

Working on some really great features over the holidays - full automation/remote control of exhaust fans and switched PDUs for the [@BearBox_io](#).

Stay tuned!



1



4



You Retweeted



Austin Storms @austorms · Nov 17, 2018

...

Also, first batch of [@BearBox_io](#) shirts came in - second batch has technical drawings on the back.

Shoot me a DM if you want one!



2



9




You Retweeted



Austin Storms @austorms · Nov 17, 2018

Weather's nice again, so I'm outside doing a little Saturday afternoon troubleshooting/component replacement. @BearBox_io



2

1

9



Show this thread

You Retweeted



Austin Storms @austorms · Nov 15, 2018

Replying to @ODELL and @MATT_ODELL



A centralized chain is worth less than worthless - it's a slow, inefficient database that you must pay to use.

0

8

32

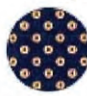


← **BearBox** 
58 Tweets
BearBox  @BearBox_io · Nov 14, 2018
Up late designing circuits for temp/humidity sensor/relay microcontrollers and having some fun with the LCD 😊



  1  9  

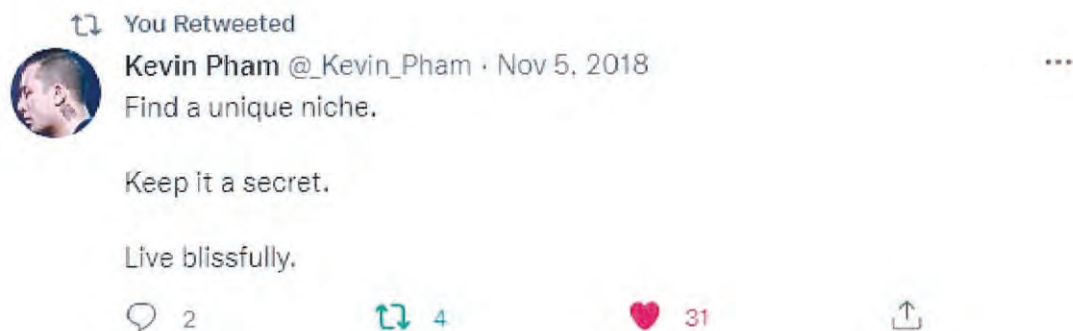
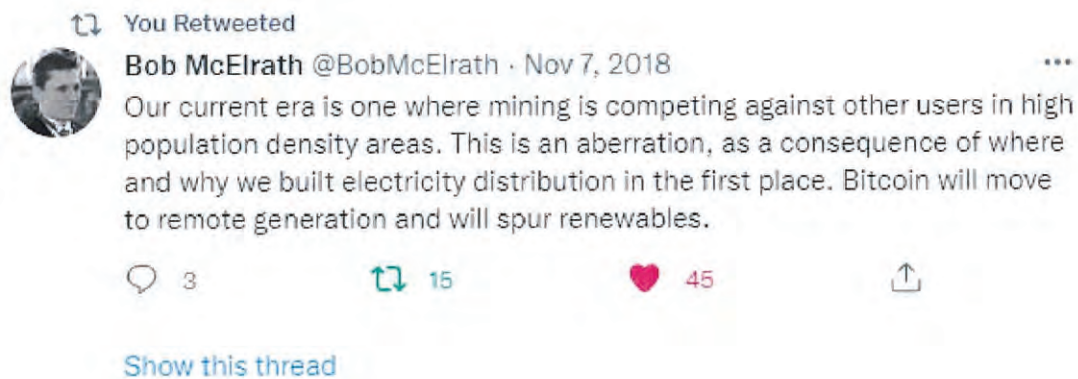
 You Retweeted

 **Bitcoin is Saving** @BitcoinIsSaving · Nov 10, 2018
The financial system that is emerging on top of the Bitcoin system is very different from the one that emerged on top of the fiat system. For example, people use BTC as collateral to borrow fiat-denominated, so BTC going up in value ("deflation") benefits both borrower and lender

 8  23  145 

[Show this thread](#)

 You Retweeted



← BearBox 
58 Tweets

 You Retweeted



Austin Storms @austorms · Nov 5, 2018

It was the perfect day to finish the @BearBox_io...

And then I bent those shutters. 🤔




 4

 4

 36



BearBox  @BearBox_io · Nov 4, 2018



← **BearBox** 
58 Tweets


  1  4  



 You Retweeted

 **Bitcoin is Saving** @BitcoinIsSaving · Nov 4, 2018 ...

Why is the Bitcoin community perpetually positive? What is it about sound money, permissionless innovation, p2p network governance that brings out the best in people? Whatever it is, I love it! As Bitcoin closes its first decade, here's to the next decade 🍷🍷🍷🍷

 17  99  591 

 You Retweeted



 **BearBox**  @BearBox_io · Nov 4, 2018 ...

Up late designing and building PDU v2 for the BearBox.

These things are going to be BADASS.





 1  4  18  

 **BearBox**  @BearBox_io · Nov 4, 2018 ...

Up late designing and building PDU v2 for the BearBox.

These things are going to be BADASS.



← **BearBox** 
58 Tweets

 **BearBox**  @BearBox_io · Nov 2, 2018
Because we like "hard money" with reliable mechanisms for restricting supply growth, we accept BTC and USD as payment for the BearBox - and USD wire transfers are assessed at a 3% premium.

Happy Friday, freaks!

 1  1  10  

 You Retweeted

 **Vortex** @theonevortex · Nov 2, 2018
"#Bitcoin  is the apex predator of money." - @danheld

 2  11  55 

 **BearBox**  @BearBox_io · Nov 1, 2018
It's Bitcoin, not blockchain.

  1  23  

 You Retweeted

 **Zack Voell** @zackvoell · Nov 1, 2018
Follow @bearbox_io and @austorms whose building efforts are par none.




 7  10  37 

BearBox @BearBox_io · Nov 1, 2018
Update: lost power around 2:45am, full restoration a few minutes later with a 14% drop in total hashrate from ~2.5PH/s to 2.15PH/s.

Manually rebooted this morning - back to normal ops (still have ~40 control boards to replace, sourced from Innosilicon).

Austin Storms @austorms · Nov 1, 2018
Currently field testing the newly installed weather shield in it's lowest position. @bearbox_io



1 1 6

BearBox @BearBox_io · Nov 1, 2018
Here's the full activity and graph from Slushpool.




Exhibit 12

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property
Organization
International Bureau

(43) International Publication Date
18 July 2019 (18.07.2019)



(10) International Publication Number
WO 2019/139632 A1

(51) International Patent Classification:

G06F 1/26 (2006.01) G06F 1/30 (2006.01)
G06F 1/20 (2006.01) G06F 1/32 (2006.01)

(21) International Application Number:

PCT/US20 18/017950

(22) International Filing Date:

13 February 2018 (13.02.2018)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

62/616,348 11 January 2018 (11.01.2018) US

(71) Applicant: LANCIUM LLC [US/US]; c/o Angelo Mikeska PLLC, 21 Waterway Avenue, Suite 300, The Woodlands, TX 77380 (US).

(72) Inventors: HENSON, David; Lancium LLC, c/o Angelo Mikeska PLLC, 21 Waterway Avenue, Suite 300, The Woodlands, TX 77380 (US). MCNAMARA, Michael; Lancium LLC, c/o Angelo Mikeska PLLC, 21 Waterway Avenue, Suite 300, The Woodlands, TX 77380 (US). CLINE, Raymond; Lancium LLC, c/o Angelo Mikeska PLLC, 21 Waterway Avenue, Suite 300, The Woodlands, TX 77380 (US).

(74) Agent: KORENCHAN, James, L.; McDonnell Boehnen Hulbert & Berghoff LLP, 300 South Wacker Drive, Suite 3200, Chicago, IL 60606 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME,

(54) Title: METHOD AND SYSTEM FOR DYNAMIC POWER DELIVERY TO A FLEXIBLE DATACENTER USING UNUTILIZED ENERGY SOURCES




FIG. 6

(57) Abstract: A flexible datacenter includes a mobile container, a behind-the-meter power input system, a power distribution system, a datacenter control system, a plurality of computing systems, and a climate control system. The datacenter control system modulates power delivery to the plurality of computing systems based on unutilized behind-the-meter power availability or an operational directive. A method of dynamic power delivery to a flexible datacenter using unutilized behind-the-meter power includes monitoring unutilized behind-the-meter power availability, determining when a datacenter ramp-up condition is met, enabling behind-the-meter power delivery to one or more computing systems when the datacenter ramp-up condition is met, and directing the one or more computing systems to perform predetermined computational operations.

[Continued on next page]

WO 2019/139632 A1

LANCIUM00000050

WO 2019/139632 A1 

MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ,
OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA,
SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN,
TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

- (84) **Designated States** (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:

— with international search report (Art. 21(3))

METHOD AND SYSTEM FOR DYNAMIC POWER DELIVERY TO
A FLEXIBLE DATACENTER USING UNUTILIZED ENERGY SOURCES

BACKGROUND OF THE INVENTION

[0001] Blockchain technology was originally conceived of as an open and distributed system for securely conducting transactions with cryptographic currency. However, the foundational principle of blockchain technology is the ability to securely transact information of any type or kind between anonymous parties without intermediaries or a centralized trust authority. As such, blockchain technology finds application outside the realm of cryptocurrency and is widely considered one of the more robust and secure means of transacting information in the computer sciences.

[0002] In typical blockchain implementations, each participating party creates a digital identity, or wallet, which includes a pair of cryptographic keys used to transact information securely and anonymously with the blockchain. The blockchain may be thought of as a constantly growing database of all prior transaction information that is securely and coherently replicated across all nodes of a peer-to-peer blockchain network. The blockchain includes a sequence of blocks, where each block includes a bundle of transactions and other data including a hash of the prior block in the chain. As such, each block in the blockchain is mathematically related to the prior block, protecting the integrity of the blockchain from the most recently added block to the genesis block in the chain. Because anyone may participate in the curation of the blockchain, once a block is added, it becomes a permanent and immutable part of the blockchain. Thus, the blockchain stores transactions in a manner that prevents the transactions from being altered or otherwise corrupted, unless all subsequent blocks in the blockchain are also altered. The immutability of the blockchain makes the malicious alteration of a block exceptionally difficult, if not impossible, and at the very least makes it easy to detect and deter any such attempt before being accepted and replicated across the blockchain network.

[0003] Each transacting party of the blockchain uses a pair of cryptographic keys to anonymously transact information. The private key is a random number maintained in secrecy by the party holder that is used to derive a public key and sign information. The private key and the public key are mathematically related such that anyone holding the public key may verify that information signed with the private key originated from the holder of the private key. When an initiating party wishes to

transact information, the information is signed with the initiating party's private key and broadcast to the blockchain network. A blockchain miner uses the initiating party's public key to verify that the initiating party initiated, or signed, the transaction. Once the initiating party's signature is validated, the transaction is validated, added to the next block in the blockchain, and replicated across all nodes.

[0004] The computational overhead of the blockchain is largely due to hashing functions used by blockchain miners to discover new blocks. While computationally intensive, the work performed by miners is critically important to the functionality of the blockchain. When an initiating party's transaction request has been lodged and the signatures validated, the transaction request is pooled in the blockchain network. Blockchain miners validate transactions and compete to discover a new block to be added to the blockchain. In order to add a newly discovered block to the blockchain, the blockchain miner must provide a cryptographic proof of the discovered block. To create the proof, the miner inputs the hash value of the prior block in the blockchain, the candidate block to be added, and a random number, commonly referred to as the nonce, to a hash function. The hash function takes input of any length and outputs an alphanumeric string of fixed length, commonly referred to as a hash, which uniquely identifies the input. However, the blockchain algorithm requires that the hash start with a certain number of leading zeros as determined by the current level of prescribed difficulty. The blockchain network modulates the level of difficulty for block discovery, by varying the number of leading zeros required in the calculated hash, based on the amount of computing power in the blockchain network.

[0005] As more computational capacity has come online, the hash rate has increased dramatically. In an effort to keep the block discovery time constant, the blockchain network modulates difficulty every 2016 blocks discovered. If the blockchain network hash rate is too high and the amount of time taken to discover a new block is less than 10 minutes, the difficulty is increased proportionally to increase the block discovery time to 10 minutes. Similarly, if the blockchain hash rate is too low and the amount of time taken to discover a new block is more than 10 minutes, the difficulty is increased proportionally to reduce the block discovery time to 10 minutes. Because there is no way to predict what hash value a given set of input data will generate, miners often have to execute the hash function a substantial number of times, each time inputting a new nonce, to generate a new hash value.

When a miner is the first to obtain a hash value having the correct number of leading zeros, they broadcast the newly discovered block to the blockchain network and the blockchain is replicated across all nodes

BRIEF SUMMARY OF THE INVENTION

- [0006] According to one aspect of one or more embodiments of the present invention, a flexible datacenter includes a mobile container, a **behind-the-meter** power input system, a power distribution system, a datacenter control system, a plurality of computing systems, and a climate control system. The datacenter control system modulates power delivery to the plurality of computing systems based on unutilized behind-the-meter power availability or an operational directive
- [0007] According to one aspect of one or more embodiments of the present invention, a method of dynamic power delivery to a flexible datacenter using unutilized behind-the-meter power includes monitoring unutilized behind-the-meter power availability, determining **when** a datacenter ramp-up condition is met, enabling behind-the-meter power delivery to one or more computing systems when the datacenter ramp-up condition is met, and directing the one or more computing systems to perform predetermined computational operations.
- [0008] Other aspects of the present invention **will** be apparent from the following description and claims

BRIEF DESCRIPTION OF THE DRAWINGS

- [0009] Figure **1** shows a computing system **in** accordance with one or more embodiments of the present invention.
- [0010] Figure **2** shows a flexible datacenter in accordance with one or more embodiments of the present invention.
- [0011] Figure **3** shows a three-phase power distribution of a flexible datacenter in accordance **with** one or more embodiments of the present invention
- [0012] Figure **4** shows a control distribution scheme of a flexible datacenter in accordance with one or more embodiments of the present invention.
- [0013] Figure **5** shows a control distribution scheme of a fleet of flexible datacenters in accordance with one or more embodiments of the present invention
- [0014] Figure **6** show's a flexible datacenter **powered** by one or more wind turbines in accordance with one or more embodiments of the present invention.

[0015] Figure 7 shows a flexible datacenter powered by one or more solar panels in accordance with one or more embodiments of the present invention.

[0016] Figure 8 shows a flexible datacenter powered by flare gas in accordance with one or more embodiments of the present invention.

[0017] Figure 9 shows a method of dynamic power delivery to a flexible datacenter using unutilized behind-the-meter power in accordance with one or more embodiments of the present invention.

DETAILED DESCRIPTION OF THE INVENTION

[0018] One or more embodiments of the present invention are described in detail with reference to the accompanying figures. For consistency, like elements in the various figures are denoted by like reference numerals. In the following detailed description of the present invention, specific details are set forth in order to provide a thorough understanding of the present invention. In other instances, well-known features to one having ordinary skill in the art are not described to avoid obscuring the description of the present invention.

[0019] Blockchain miners are typically compensated for their efforts through either a discover}, fee or a fee paid by one or more of the transacting parties. Consequently, more and more computing resources are coming online to compete for these fees. As the number of computing resources increases, the blockchain network modulates the difficulty level, requiring hash values with more leading zeros. In essence, the increased difficulty means more hashing operations are required to find a valid hash. As such, there is an increasing number of computing resources executing an increasing number of hash functions that do not result in the discovery of a valid hash, yet still consume a substantial amount of power.

[0020] The intensive computational demand of blockchain applications makes the widespread adoption of blockchain technology inefficient and unsustainable from an energy and environmental perspective. In certain blockchain applications, with limited participation, roughly 5 quintillion 256-bit cryptographic hashes are created each and every' second of every day. While it is difficult to determine how much energy is required for that computational task, it is estimated to be in excess of 500 megawatts, the vast majority of which is sourced from fossil fuels. The majority of blockchain mining operations are currently being conducted in the People's Republic of China and powered by coal-fired energy. As blockchain technology

proliferates, there is concern that the energy required to sustain such blockchain applications could exceed that of a developed country.

[0021] While future versions of blockchain technology may improve power consumption for various blockchain operations, including hashing functions, industry' efforts have focused on the development of central processing units ("CPUs"), graphics processing units ("GPUs"), and application specific integrated circuits ("ASICs") that are specifically designed to perform blockchain operations in a more efficient manner. While such efforts are beneficial, the issue remains, the widespread adoption of blockchain technology will require substantially more power than is economically and environmentally feasible.

[0022] Accordingly, in one or more embodiments of the present invention, a method and system for dynamic power delivery' to a flexible datacenter uses unutilized behind-the-meter power sources without transmission and distribution costs. The flexible datacenter may be configured to modulate power delivery to one or more computing systems based on the availability of unutilized behind-the-meter power or an operational directive. For example, the flexible datacenter may ramp-up to a fully online status, ramp-down to a fully offline status, or dynamically reduce power consumption, act a load balancer, or adjust the power factor. Advantageously, the flexible datacenter may perform computational operations, such as blockchain hashing operations, with little to no energy costs, using clean and renewable energy that would otherwise be wasted.

[0023] Figure 1 shows a computing system 100 in accordance with one or more embodiments of the present invention. Computing system 100 may include one or more central processing units (singular "CPU" or plural "CPUs") 105, host bridge 110, input/output ("IO") bridge 115, graphics processing units (singular "GPU" or plural "GPUs") 125, and/or application-specific integrated circuits (singular "ASIC" or plural "ASICs") (not shown) disposed on one or more printed circuit boards (not shown) that are configured to perform computational operations. Each of the one or more CPUs 105, GPUs 125, or ASICs (not shown) may be a single-core (not independently illustrated) device or a multi-core (not independently illustrated) device. Multi-core devices typically include a plurality of cores (not shown) disposed on the same physical die (not shown) or a plurality of cores (not shown) disposed on multiple die (not shown) that are collectively disposed within the same mechanical package (not shown).

[0024] CPU 105 may be a general purpose computational device typically configured to execute software instructions. CPU 105 may include an interface 108 to host bridge 110, an interface 118 to system memory 120, and an interface 123 to one or more IO devices, such as, for example, one or more GPUs 125. GPU 125 may serve as a specialized computational device typically configured to perform graphics functions related to frame buffer manipulation. However, one of ordinary skill in the art will recognize that GPU 125 may be used to perform non-graphics related functions that are computationally intensive. In certain embodiments, GPU 125 may interface 123 directly with CPU 125 (and interface 118 with system memory 120 through CPU 105). In other embodiments, GPU 125 may interface 121 with host bridge 110 (and interlace 116 or 118 with system memory 120 through host bridge 110 or CPU 105 depending on the application or design). In still other embodiments, GPU 125 may interface 133 with IO bridge 115 (and interface 116 or 118 with system memory 120 through host bridge 110 or CPU 105 depending on the application or design). The functionality of GPU 125 may be integrated, in whole or in part, with CPU 105.

[0025] Host bridge 110 may be an interface device configured to interface between the one or more computational devices and IO bridge 115 and, in some embodiments, system memory 120. Host bridge 110 may include an interface 108 to CPU 105, an interface 113 to IO bridge 115, for embodiments where CPU 105 does not include an interface 118 to system memory 120, an interface 116 to system memory 120, and for embodiments where CPU 105 does not include an integrated GPU 125 or an interface 123 to GPU 125, an interface 121 to GPU 125. The functionality of host bridge 110 may be integrated, in whole or in part, with CPU 105. IO bridge 115 may be an interface device configured to interface between the one or more computational devices and various IO devices (*e.g.*, 140, 145) and IO expansion, or add-on, devices (not independently illustrated). IO bridge 115 may include an interface 113 to host bridge 110, one or more interfaces 133 to one or more IO expansion devices 135, an interface 138 to keyboard 140, an interface 143 to mouse 145, an interface 148 to one or more local storage devices 150, and an interface 153 to one or more network interface devices 155. The functionality of IO bridge 115 may be integrated, in whole or in part, with CPU 105 or host bridge 110. Each local storage device 150, if any, may be a solid-state memory device, a solid-state memory device array, a hard disk drive, a hard disk drive array, or any other

non-transitory computer readable medium. Network interface device 155 may provide one or more network interfaces including any network protocol suitable to facilitate networked communications.

[0026] Computing system 100 may include one or more network-attached storage devices 160 in addition to, or instead of, one or more local storage devices 150. Each network-attached storage device 160, if any, may be a solid-state memory device, a solid-state memory device array, a hard disk drive, a hard disk drive array, or any other **non-transitory** computer readable medium. **Network-attached** storage device 160 may or may **not** be collocated **with** computing system 100 and may be accessible to computing system 100 via one or more **network** interfaces provided by one or more **network** interface devices 155.

[0027] One of ordinary skill in the art **will** recognize that computing system 100 may be a conventional computing system or an application-specific computing system. **In** certain embodiments, an application-specific computing system may include one or more ASICs (not shown) that are configured to perform one or more functions, such as hashing, in a more efficient manner. The one or more ASICs (not shown) may interface directly with CPU 105, host bridge 110, or GPU 125 or interface through **IO** bridge **115**. Alternatively, **in** other embodiments, an application-specific computing system may be reduced to only those components **necessary** to perform a desired function in an effort to reduce one or more of chip count, printed circuit board footprint, thermal design power, and power consumption. The one or more ASICs (not shown) may be used instead of one or more of CPU 105, host bridge **110**, **IO** bridge **115**, or GPU 125. **In** such systems, the one or more ASICs may incorporate sufficient functionality to perform certain network and computational functions in a minimal footprint **with** substantially fewer component devices.

[0028] As such, one of **ordinary** skill in the art will recognize that CPU 105, host bridge **110**, **IO** bridge **115**, GPU 125, or ASIC (not shown) or a subset, superset, or combination of functions or features thereof, may be integrated, distributed, or excluded, in whole or in part, based on an application, design, or **form** factor in accordance with one or more embodiments of the present invention. Thus, the description of computing system 100 is merely exemplary and not intended to **limit** the **type**, kind, or configuration of component devices that constitute a computing system 100 suitable for performing computing operations in accordance with one or more embodiments of the present invention.

- [0029] One of ordinary skill in the art will recognize that computing system 100 may be a stand alone, laptop, desktop, server, blade, or rack mountable system and may vary based on an application or design.
- [0030] Figure 2 shows a flexible datacenter 200 in accordance with one or more embodiments of the present invention. Flexible datacenter 200 may include a mobile container 205, a behind-the-meter power input system 210, a power distribution system 215, a climate control system (*e.g.*, 250, 260, 270, 280, and/or 290), a datacenter control system 220, and a plurality of computing systems 100 disposed in one or more racks 240. Datacenter control system 220 may be a computing system (*e.g.*, 100 of Figure 1) configured to dynamically modulate power delivery to one or more computing systems 100 disposed within flexible datacenter 200 based on unutilized behind-the-meter power availability or an operational directive from a local station control system (not shown), a remote master control system (not shown), or a grid operator (not shown).
- [0031] In certain embodiments, mobile container 205 may be a storage trailer disposed on wheels and configured for rapid deployment. In other embodiments, mobile container 205 may be a storage container (not shown) configured for placement on the ground and potentially stacked in a vertical manner (not shown). In still other embodiments, mobile container 205 may be an inflatable container, a floating container, or any other type or kind of container suitable for housing a mobile datacenter 200.
- [0032] Flexible datacenter 200 may be rapidly deployed on site near a source of unutilized behind-the-meter power generation. Behind-the-meter power input system 210 may be configured to input power to flexible datacenter 200. Behind-the-meter power input system 210 may include a first input (not independently illustrated) configured to receive three-phase behind-the-meter alternating current (“AC”) voltage. In certain embodiments, behind-the-meter power input system 210 may include a supervisory AC-to-AC step-down transformer (not shown) configured to step down three-phase behind-the-meter AC voltage to single-phase supervisory nominal AC voltage or a second input (not independently illustrated) configured to receive single-phase supervisory nominal AC voltage from the local station (not shown) or a metered source (not shown). Behind-the-meter power input system 210 may provide single-phase supervisory nominal AC voltage to datacenter control system 220, which may remain powered at almost all times to control the

operation of flexible datacenter 200. The first input (not independently illustrated) or a third input (not independently illustrated) of behind-the-meter power input system 210 may direct three-phase behind-the-meter AC voltage to an operational AC-to-AC step-down transformer (not shown) configured to controllably step down three-phase behind-the-meter AC voltage to three-phase nominal AC voltage. Datacenter control system 220 may controllably enable or disable generation or provision of three-phase nominal AC voltage by the operational AC-to-AC step-down transformer (not shown).

[0033] Behind-the-meter power input system 210 may provide three phases of three-phase nominal AC voltage to power distribution system 215. Power distribution system 215 may controllably provide a single phase of three-phase nominal AC voltage to each computing system 100 or group 240 of computing systems 100 disposed within flexible datacenter 200. Datacenter control system 220 may controllably select which phase of three-phase nominal AC voltage that power distribution system 215 provides to each computing system 100 or group 240 of computing systems 100. In this way, datacenter control system 220 may modulate power delivery by either ramping-up flexible datacenter 200 to fully operational status, ramping-down flexible datacenter 200 to offline status (where only datacenter control system 220 remains powered), reducing power consumption by withdrawing power delivery from, or reducing power to, one or more computing systems 100 or groups 240 of computing systems 100, or modulating a power factor correction factor for the local station by controllably adjusting which phases of three-phase nominal AC voltage are used by one or more computing systems 100 or groups 240 of computing systems 100.

[0034] Flexible datacenter 200 may include a climate control system (*e.g.*, 250, 260, 270, 280, 290) configured to maintain the plurality of computing systems 100 within their operational temperature range. In certain embodiments, the climate control system may include an air intake 250, an evaporative cooling system 270, a fan 280, and an air outtake 260. In other embodiments, the climate control system may include an air intake 250, an air conditioner or refrigerant cooling system 290, and an air outtake 260. In still other embodiments, the climate control system may include a computer room air conditioner system (not shown), a computer room air handler system (not shown), or an immersive cooling system (not shown). One of ordinary skill in the art will recognize that any suitable heat extraction system (not

shown) configured to maintain the operation of the plurality of computing systems 100 within their operational temperature range may be used in accordance with one or more embodiments of the present invention.

[0035] Flexible datacenter 200 may include a battery' system (not shown) configured to convert three-phase nominal AC voltage to nominal DC voltage and store power in a plurality of storage cells. The battery' system (not shown) may include a DC-to-AC inverter configured to convert nominal DC voltage to three-phase nominal AC voltage for flexible datacenter 200 use. Alternatively, the battery system (not shown) may include a DC-to-AC inverter configured to convert nominal DC voltage to single-phase nominal AC voltage to power datacenter control system 220.

[0036] One of ordinary skill in the art will recognize that a voltage level of three-phase behind-the-meter AC voltage may vary based on an application or design and the type or kind of local power generation. As such, a type, kind, or configuration of the operational AC-to-AC step down transformer (not shown) may vary based on the application or design. In addition, the frequency and voltage level of three-phase nominal AC voltage, single-phase nominal AC voltage, and nominal DC voltage may vary based on the application or design in accordance with one or more embodiments of the present invention.

[0037] Figure 3 shows a three-phase power distribution of a flexible datacenter 200 in accordance with one or more embodiments of the present invention. Flexible datacenter 200 may include a plurality of racks 240, each of which may include one or more computing systems 100 disposed therein. As discussed above, the behind-the-meter power input system (210 of Figure 2) may provide three phases of three-phase nominal AC voltage to the power distribution system (215 of Figure 2). The power distribution system (215 of Figure 2) may controllably provide a single phase of three-phase nominal AC voltage to each computing system 100 or group 240 of computing systems 100 disposed within flexible datacenter 200. For example, a flexible datacenter 200 may include eighteen racks 240, each of which may include eighteen computing systems 100. The power distribution system (215 of Figure 2) may control which phase of three-phase nominal AC voltage is provided to one or more computing systems 100, a rack 240 of computing systems 100, or a group (*e.g.*, 310, 320, or 330) of racks 240 of computing systems 100.

[0038] In the figure, for purposes of illustration only, eighteen racks 240 are divided into a first group of six racks 310, a second group of six racks 320, and a third group

of six racks 330, where each rack contains eighteen computing systems 100. The power distribution system (215 of Figure 2) may, for example, provide a first phase of three-phase nominal AC voltage to the first group of six racks 310, a second phase of three-phase nominal AC voltage to the second group of six racks 320, and a third phase of three-phase nominal AC voltage to the third group of six racks 330. If the flexible datacenter (200 of Figure 2) receives an operational directive from the local station (not shown) to provide power factor correction, the datacenter control system (220 of Figure 2) may direct the power distribution system (215 of Figure 2) to adjust which phase or phases of three-phase nominal AC voltage are used to provide the power factor correction required by the local station (not shown) or grid operator (not shown). One of ordinary skill in the art will recognize that, in addition to the power distribution, the load may be varied by adjusting the number of computing systems 100 operatively powered. As such, the flexible datacenter (200 of Figure 2) may be configured to act as a capacitive or inductive load to provide the appropriate reactance necessary to achieve the power factor correction required by the local station (not shown).

[0039] Figure 4 shows a control distribution scheme of a flexible datacenter 200 in accordance with one or more embodiments of the present invention. Datacenter control system 220 may independently, or cooperatively with one or more of local station control system 410, remote master control system 420, and grid operator 440, modulate power delivery to flexible datacenter 200. Specifically, power delivery may be dynamically adjusted based on conditions or operational directives.

[0040] Local station control system 410 may be a computing system (*e.g.*, 100 of Figure 1) that is configured to control various aspects of the local station (not independently illustrated) that generates power and sometimes generates unutilized behind-the-meter power. Local station control system 410 may communicate with remote master control system 420 over a networked connection 430 and with datacenter control system 220 over a networked or hardwired connection 415. Remote master control system 420 may be a computing system (*e.g.*, 100 of Figure 1) that is located offsite, but connected via a network connection 425 to datacenter control system 220, that is configured to provide supervisory or override control of flexible datacenter 200 or a fleet (not shown) of flexible datacenters 200. Grid operator 440 may be a computing system (*e.g.*, 100 of Figure 1) that is configured to control various aspects of the grid (not independently illustrated) that receives

power from the local station (not independently illustrated). Grid operator 440 may communicate with local station control system 440 over a networked or hardwired connection 445.

[0041] Datacenter control system 220 may monitor unutilized behind-the-meter power availability at the local station (not independently illustrated) and determine when a datacenter ramp-up condition is met. Unutilized behind-the-meter power availability may include one or more of excess local power generation, excess local power generation that the grid cannot accept, local power generation that is subject to economic curtailment, local power generation that is subject to reliability curtailment, local power generation that is subject to power factor correction, situations where local power generation is prohibitively low, start up situations, transient situations, or testing situations where there is an economic advantage to using locally generated behind-the-meter power generation, specifically power available at little to no cost and with no associated transmission or distribution costs.

[0042] The datacenter ramp-up condition may be met if there is sufficient behind-the-meter power availability and there is no operational directive from local station control system 410, remote master control system 420, or grid operator 440 to go offline or reduce power. As such, datacenter control system 220 may enable 435 behind-the-meter power input system 210 to provide three-phase nominal AC voltage to the power distribution system (215 of Figure 2) to power the plurality of computing systems (100 of Figure 2) or a subset thereof. Datacenter control system 220 may optionally direct one or more computing systems (100 of Figure 2) to perform predetermined computational operations. For example, if the one or more computing systems (100 of Figure 2) are configured to perform blockchain hashing operations, datacenter control system 220 may direct them to perform blockchain hashing operations for a specific blockchain application, such as, for example, Bitcoin, Litecoin, or Ethereum. Alternatively, one or more computing systems (100 of Figure 2) may be configured to independently receive a computational directive from a network connection (not shown) to a peer-to-peer blockchain network (not shown) such as, for example, a network for a specific blockchain application, to perform predetermined computational operations.

[0043] Remote master control system 420 may specify to datacenter control system 220 what sufficient behind-the-meter power availability constitutes, or datacenter

control system 220 may be programmed with a predetermined preference or criteria on which to make the determination independently. For example, in certain circumstances, sufficient behind-the-meter power availability may be less than that required to fully power the entire flexible datacenter 200. In such circumstances, datacenter control system 220 may provide power to only a subset of computing systems (100 of Figure 2), or operate the plurality of computing systems (100 of Figure 2) in a lower power mode, that is within the sufficient, but less than full, range of power that is available.

[0044] While flexible datacenter 200 is online and operational, a datacenter ramp-down condition may be met when there is insufficient, or anticipated to be insufficient, behind-the-meter power availability or there is an operational directive from local station control system 410, remote master control system 420, or grid operator 440. Datacenter control system 220 may monitor and determine when there is insufficient, or anticipated to be insufficient, behind-the-meter power availability. As noted above, sufficiency may be specified by remote master control system 420 or datacenter control system 220 may be programmed with a predetermined preference or criteria on which to make the determination independently. An operational directive may be based on current dispatchability, forward looking forecasts for when unutilized behind-the-meter power is, or is expected to be, available, economic considerations, reliability considerations, operational considerations, or the discretion of the local station 410, remote master control 420, or grid operator 440. For example, local station control system 410, remote master control system 420, or grid operator 440 may issue an operational directive to flexible datacenter 200 to go offline and power down. When the datacenter ramp-down condition is met, datacenter control system 220 may disable power delivery to the plurality of computing systems (100 of Figure 2). Datacenter control system 220 may disable 435 behind-the-meter power input system 210 from providing three-phase nominal AC voltage to the power distribution system (215 of Figure 2) to power down the plurality of computing systems (100 of Figure 2), while datacenter control system 220 remains powered and is capable of rebooting flexible datacenter 200 when unutilized behind-the-meter power becomes available again.

[0045] While flexible datacenter 200 is online and operational, changed conditions or an operational directive may cause datacenter control system 220 to modulate power consumption by flexible datacenter 200. Datacenter control system 220 may

determine, or local station control system 410, remote master control system 420, or grid operator 440 may communicate, that a change in local conditions may result in less power generation, availability, or economic feasibility, than would be necessary to fully power flexible datacenter 200. In such situations, datacenter control system 220 may take steps to reduce or stop power consumption by flexible datacenter 200 (other than that required to maintain operation of datacenter control system 220). Alternatively, local station control system 410, remote master control system 420, or grid operator 440, may issue an operational directive to reduce power consumption for any reason, the cause of which may be unknown. In response, datacenter control system 220 may dynamically reduce or withdraw power delivery to one or more computing systems (100 of Figure 2) to meet the dictate. Datacenter control system 220 may controllably provide three-phase nominal AC voltage to a smaller subset of computing systems (100 of Figure 2) to reduce power consumption. Datacenter control system 220 may dynamically reduce the power consumption of one or more computing systems (100 of Figure 2) by reducing their operating frequency or forcing them into a lower power mode through a network directive.

[0046] One of ordinary skill in the art will recognize that datacenter control system 220 may be configured to have a number of different configurations, such as a number or type or kind of computing systems (100 of Figure 2) that may be powered, and in what operating mode, that correspond to a number of different ranges of sufficient and available unutilized behind-the-meter power availability. As such, datacenter control system 220 may modulate power delivery over a variety of ranges of sufficient and available unutilized behind-the-meter power availability.

[0047] Figure 5 shows a control distribution of a fleet 500 of flexible datacenters 200 in accordance with one or more embodiments of the present invention. The control distribution of a flexible datacenter 200 shown and described with respect to Figure 4 may be extended to a fleet 500 of flexible datacenters 200. For example, a first local station (not independently illustrated), such as, for example, a wind farm (not shown), may include a first plurality 510 of flexible datacenters 200a through 200d, which may be collocated or distributed across the local station (not shown). A second local station (not independently illustrated), such as, for example, another wind farm or a solar farm (not shown), may include a second plurality 520 of flexible datacenters 200e through 200y which may be collocated or distributed

across the local station (not shown). One of ordinary skill in the art will recognize that the number of flexible datacenters 200 deployed at a given station and the number of stations within the fleet may vary based on an application or design in accordance with one or more embodiments of the present invention.

[0048] Remote master control system 420 may provide supervisory control over fleet 500 of flexible datacenters 200 in a similar manner to that shown and described with respect to Figure 4, with the added flexibility to make high level decisions with respect to fleet 500 that may be counterintuitive to a given station. Remote master control system 420 may make decisions regarding the issuance of operational directives to a given local station based on, for example, the status of each local station where flexible datacenters 200 are deployed, the workload distributed across fleet 500, and the expected computational demand required for the expected workload. In addition, remote master control system 420 may shift workloads from a first plurality 510 of flexible datacenters 200 to a second plurality 520 of flexible datacenters 200 for any reason, including, for example, a loss of unutilized behind-the-meter power availability at one local station and the availability of unutilized behind-the-meter power at another local station.

[0049] Figure 6 shows a flexible datacenter 200 powered by one or more wind turbines 610 in accordance with one or more embodiments of the present invention. A wind farm 600 typically includes a plurality of wind turbines 610, each of which intermittently generates a wind-generated AC voltage. The wind-generated AC voltage may vary based on a type, kind, or configuration of farm 600, turbine 610, and incident wind speed. The wind-generated AC voltage is typically input into a turbine AC-to-AC step-up transformer (not shown) that is disposed within the nacelle (not independently illustrated) or at the base of the mast (not independently illustrated) of turbine 610. The turbine AC-to-AC step up transformer (not shown) outputs three-phase wind-generated AC voltage 620. Three-phase wind-generated AC voltage 620 produced by the plurality of wind turbines 610 is collected 625 and provided 630 to another AC-to-AC step-up transformer 640 that steps up three-phase wind-generated AC voltage 620 to three-phase grid AC voltage 650 suitable for delivery to grid 660. Three-phase grid AC voltage 650 may be stepped down with an AC-to-AC step-down transformer 670 configured to produce three-phase local station AC voltage 680 provided to local station 690. One of ordinary skill in the art will recognize that the actual voltage levels may vary based on the type, kind,

or number of wind turbines 610, the configuration or design of wind farm 600, and grid 660 that it feeds into.

[0050] The output side of AC-to-AC step-up transformer 640 that connects to grid 660 may be metered and is typically subject to transmission and distribution costs. In contrast, power consumed on the input side of AC-to-AC step-up transformer 640 may be considered behind-the-meter and is typically not subject to transmission and distribution costs. As such, one or more flexible datacenters 200 may be powered by three-phase wind-generated AC voltage 620. Specifically, in wind farm 600 applications, the three-phase behind-the-meter AC voltage used to power flexible datacenter 200 may be three-phase wind-generated AC voltage 620. As such, flexible datacenter 200 may reside behind-the-meter, avoid transmission and distribution costs, and may be dynamically powered when unutilized behind-the-meter power is available.

[0051] Unutilized behind-the-meter power availability may occur when there is excess local power generation. In high wind conditions, wind farm 600 may generate more power than, for example, AC-to-AC step-up transformer 640 is rated for. In such situations, wind farm 600 may have to take steps to protect its equipment from damage, which may include taking one or more turbines 610 offline or shunting their voltage to dummy loads or ground. Advantageously, one or more flexible datacenters 200 may be used to consume power on the input side of AC-to-AC step-up transformer 640, thereby allowing wind farm 600 to operate equipment within operating ranges while flexible datacenter 200 receives behind-the-meter power without transmission or distribution costs. The local station control system (not independently illustrated) of local station 690 may issue an operational directive to the one or more flexible datacenters 200 or to the remote master control system (420 of Figure 4) to ramp-up to the desired power consumption level. When the operational directive requires the cooperative action of multiple flexible datacenters 200, the remote master control system (420 of Figure 4) may determine how to power each individual flexible datacenter 200 in accordance with the operational directive or provide an override to each flexible datacenter 200.

[0052] Another example of unutilized behind-the-meter power availability is when grid 660 cannot, for whatever reason, take the power being produced by wind farm 600. In such situations, wind farm 600 may have to take one or more turbines 610 offline or shunt their voltage to dummy loads or ground. Advantageously, one or

more flexible datacenters 200 may be used to consume power on the input side of AC-to-AC step-up transformer 640, thereby allowing wind farm 600 to either produce power to grid 660 at a lower level or shut down transformer 640 entirely while flexible datacenter 200 receives behind-the-meter power without transmission or distribution costs. The local station control system (not independently illustrated) of local station 690 or the grid operator (not independently illustrated) of grid 660 may issue an operational directive to the one or more flexible datacenters 200 or to the remote master control system (420 of Figure 4) to ramp-up to the desired power consumption level. When the operational directive requires the cooperative action of multiple flexible datacenters 200, the remote master control system (420 of Figure 4) may determine how to power each individual flexible datacenter 200 in accordance with the operational directive or provide an override to each flexible datacenter 200.

[0053] Another example of unutilized behind-the-meter power availability is when wind farm 600 is selling power to grid 660 at a negative price that is offset by a production tax credit. In certain circumstances, the value of the production tax credit may exceed the price wind farm 600 would have to pay to grid 660 to offload their generated power. Advantageously, one or more flexible datacenters 200 may be used to consume power behind-the-meter, thereby allowing wind farm 600 to produce and obtain the production tax credit, but sell less power to grid 660 at the negative price. The local station control system (not independently illustrated) of local station 690 may issue an operational directive to the one or more flexible datacenters 200 or to the remote master control system (420 of Figure 4) to ramp-up to the desired power consumption level. When the operational directive requires the cooperative action of multiple flexible datacenter 200, the remote master control system (420 of Figure 4) may determine how to power each individual flexible datacenter 200 in accordance with the operational directive or provide an override to each flexible datacenter 200.

[0054] Another example of unutilized behind-the-meter power availability is when wind farm 600 is selling power to grid 660 at a negative price because grid 660 is oversupplied or is instructed to stand down and stop producing altogether. The grid operator (not independently illustrated) may select certain power generation stations to go offline and stop producing power to grid 660. Advantageously, one or more flexible datacenters 200 may be used to consume power behind-the-meter, thereby

allowing wind farm 600 to stop producing power to grid 660, but making productive use of the power generated behind-the-meter without transmission or distribution costs. The local station control system (not independently illustrated) of the local station 690 or the grid operator (not independently illustrated) of grid 660 may issue an operational directive to the one or more flexible datacenters 200 or to the remote master control system (420 of Figure 4) to ramp-up to the desired power consumption level. When the operational directive requires the cooperative action of multiple flexible datacenters 200, the remote master control system (420 of Figure 4) may determine how to power each individual flexible datacenter 200 in accordance with the operational directive or provide an override to each flexible datacenter 200.

[0055] Another example of unutilized behind-the-meter power availability is when wind farm 600 is producing power to grid 660 that is unstable, out of phase, or at the wrong frequency, or grid 660 is already unstable, out of phase, or at the wrong frequency for whatever reason. The grid operator (not independently illustrated) may select certain power generation stations to go offline and stop producing power to grid 660. Advantageously, one or more flexible datacenters 200 may be used to consume power behind-the-meter, thereby allowing wind farm 600 to stop producing power to grid 660, but make productive use of the power generated behind-the-meter without transmission or distribution costs. The local station control system (not independently illustrated) of local station 690 may issue an operational directive to the one or more flexible datacenters 200 or to the remote master control system (420 of Figure 4) to ramp-up to the desired power consumption level. When the operational directive requires the cooperative action of multiple flexible datacenters 200, the remote master control system (420 of Figure 4) may determine how to power each individual flexible datacenter 200 in accordance with the operational directive or provide an override to each flexible datacenter 200.

[0056] Further examples of unutilized behind-the-meter power availability is when wind farm 600 experiences low wind conditions that make it not economically feasible to power up certain components, such as, for example, the local station (not independently illustrated), but there may be sufficient behind-the-meter power availability to power one or more flexible datacenters 200. Similarly, unutilized behind-the-meter power availability may occur when wind farm 600 is starting up,

or testing, one or more turbines 610. Turbines 610 are frequently offline for installation, maintenance, and sendee and must be tested prior to coming online as part of the array. One or more flexible datacenters 200 may be powered by one or more turbines 610 that are offline from farm 600. The above-noted examples of when unutilized behind-the-meter power is available are merely exemplary and are not intended to limit the scope of what one of ordinary skill in the art would recognize as unutilized behind-the-meter power availability. Unutilized behind-the-meter power availability may occur anytime there is power available and accessible behind-the-meter that is not subject to transmission and distribution costs and there is an economic advantage to using it.

[0057] One of ordinary skill in the art will recognize that wind farm 600 and wind turbine 610 may vary based on an application or design in accordance with one or more embodiments of the present invention.

[0058] Figure 7 shows a flexible datacenter 200 powered by one or more solar panels 710 in accordance with one or more embodiments of the present invention. A solar farm 700 typically includes a plurality of solar panels 710, each of which intermittently generates a solar-generated DC voltage 720. Solar-generated DC voltage 720 may vary based on a type, kind, or configuration of farm 700, panel 710, and incident sunlight. Solar-generated DC voltage 720 produced by the plurality of solar panels 710 is collected 725 and provided 730 to a DC-to-AC inverter that converts solar-generated DC voltage into three-phase solar-generated AC voltage 750. Three-phase solar-generated AC voltage 750 is provided to an AC-to-AC step-up transformer 760 that steps up three-phase solar-generated AC voltage to three-phase grid AC voltage 790. Three-phase grid AC voltage 790 may be stepped down with an AC-to-AC step-down transformer 785 configured to produce three-phase local station AC voltage 777 provided to local station 775. One of ordinary skill in the art will recognize that the actual voltage levels may vary' based on the type, kind, or number of solar panels 710, the configuration or design of solar farm 700, and grid 790 that it feeds into.

[0059] The output side of AC-to-AC step-up transformer 760 that connects to grid 790 may be metered and is typically subject to transmission and distribution costs. In contrast, power consumed on the input side of AC-to-AC step-up transformer 760 may be considered behind-the-meter and is typically not subject to transmission and distribution costs. As such, one or more flexible datacenters 200 may be powered by

three-phase solar-generated AC voltage 750. Specifically, in solar farm 700 applications, the three-phase behind-the-meter AC voltage used to power flexible datacenter 200 may be three-phase solar-generated AC voltage 750. As such, flexible datacenter 200 may reside behind-the-meter, avoid transmission and distribution costs, and may be dynamically powered when unutilized behind-the-meter power is available.

[0060] Unutilized behind-the-meter power availability may occur when there is excess local power generation. In high incident sunlight situations, solar farm 700 may generate more power than, for example, AC-to-AC step-up transformer 760 is rated for. In such situations, solar farm 700 may have to take steps to protect its equipment from damage, which may include taking one or more panels 710 offline or shunting their voltage to dummy loads or ground. Advantageously, one or more flexible datacenters 200 may be used to consume power on the input side of AC-to-AC step-up transformer 760, thereby allowing solar farm 700 to operate equipment within operating ranges while flexible datacenter 200 receives behind-the-meter power without transmission or distribution costs. The local station control system (not independently illustrated) of local station 775 may issue an operational directive to the one or more flexible datacenters 200 or to the remote master control system (420 of Figure 4) to ramp-up to the desired power consumption level. When the operational directive requires the cooperative action of multiple flexible datacenters 200, the remote master control system (420 of Figure 4) may determine how to power each individual flexible datacenter 200 in accordance with the operational directive or provide an override to each flexible datacenter 200.

[0061] Another example of unutilized behind-the-meter power availability is when grid 790 cannot, for whatever reason, take the power being produced by solar farm 700. In such situations, solar farm 700 may have to take one or more panels 710 offline or shunt their voltage to dummy loads or ground. Advantageously, one or more flexible datacenters 200 may be used to consume power on the input side of AC-to-AC step-up transformer 760, thereby allowing solar farm 700 to either produce power to grid 790 at a lower level or shut down transformer 760 entirely while flexible datacenter 200 receives behind-the-meter power without transmission or distribution costs. The local station control system (not independently illustrated) of local station 775 or the grid operator (not independently illustrated) of grid 790 may issue an operational directive to the one or more flexible datacenters 200 or to

the remote master control system (420 of Figure 4) to ramp-up to the desired power consumption level. When the operational directive requires the cooperative action of multiple flexible datacenters 200, the remote master control system (420 of Figure 4) may determine how to power each individual flexible datacenter 200 in accordance with the operational directive or provide an override to each flexible datacenter 200.

[0062] Another example of unutilized behind-the-meter power availability is when solar farm 700 is selling power to grid 790 at a negative price that is offset by a production tax credit. In certain circumstances, the value of the production tax credit may exceed the price solar farm 700 would have to pay to grid 790 to offload their generated power. Advantageously, one or more flexible datacenters 200 may be used to consume power behind-the-meter, thereby allowing solar farm 700 to produce and obtain the production tax credit, but sell less power to grid 790 at the negative price. The local station control system (not independently illustrated) of local station 775 may issue an operational directive to the one or more flexible datacenters 200 or to the remote master control system (420 of Figure 4) to ramp-up to the desired power consumption level. When the operational directive requires the cooperative action of multiple flexible datacenter 200, the remote master control system (420 of Figure 4) may determine how to power each individual flexible datacenter 200 in accordance with the operational directive or provide an override to each flexible datacenter 200.

[0063] Another example of unutilized behind-the-meter power availability is when solar farm 700 is selling power to grid 790 at a negative price because grid 790 is oversupplied or is instructed to stand down and stop producing altogether. The grid operator (not independently illustrated) may select certain power generation stations to go offline and stop producing power to grid 790. Advantageously, one or more flexible datacenters 200 may be used to consume power behind-the-meter, thereby allowing solar farm 700 to stop producing power to grid 790, but making productive use of the power generated behind-the-meter without transmission or distribution costs. The local station control system (not independently illustrated) of the local station 775 or the grid operator (not independently illustrated) of grid 790 may issue an operational directive to the one or more flexible datacenters 200 or to the remote master control system (420 of Figure 4) to ramp-up to the desired power consumption level. When the operational directive requires the cooperative action of

multiple flexible datacenters 200, the remote master control system (420 of Figure 4) may determine how to power each individual flexible datacenter 200 in accordance with the operational directive or provide an override to each flexible datacenter 200.

[0064] Another example of unutilized behind-the-meter power availability is when solar farm 700 is producing power to grid 790 that is unstable, out of phase, or at the wrong frequency, or grid 790 is already unstable, out of phase, or at the wrong frequency for whatever reason. The grid operator (not independently illustrated) may select certain power generation stations to go offline and stop producing power to grid 790. Advantageously, one or more flexible datacenters 200 may be used to consume power behind-the-meter, thereby allowing solar farm 700 to stop producing power to grid 790, but make productive use of the power generated behind-the-meter without transmission or distribution costs. The local station control system (not independently illustrated) of local station 775 may issue an operational directive to the one or more flexible datacenters 200 or to the remote master control system (420 of Figure 4) to ramp-up to the desired power consumption level. When the operational directive requires the cooperative action of multiple flexible datacenters 200, the remote master control system (420 of Figure 4) may determine how to power each individual flexible datacenter 200 in accordance with the operational directive or provide an override to each flexible datacenter 200.

[0065] Further examples of unutilized behind-the-meter power availability is when solar farm 700 experiences intermittent cloud cover such that it is not economically feasible to power up certain components, such as, for example local station 775, but there may be sufficient behind-the-meter power availability to power one or more flexible datacenters 200. Similarly, unutilized behind-the-meter power availability may occur when solar farm 700 is starting up, or testing, one or more panels 710. Panels 710 are frequently offline for installation, maintenance, and service and must be tested prior to coming online as part of the array. One or more flexible datacenters 200 may be powered by one or more panels 710 that are offline from farm 700. The above-noted examples of when unutilized behind-the-meter power is available are merely exemplary- and are not intended to limit the scope of what one of ordinary skill in the art would recognize as unutilized behind-the-meter power availability. Behind-the-meter power availability may occur anytime there is power

available and accessible behind-the-meter that is not subject to transmission and distribution costs and there is an economic advantage to using it.

[0066] One of ordinary skill in the art will recognize that solar farm 700 and solar panel 710 may vary based on an application or design in accordance with one or more embodiments of the present invention.

[0067] Figure 8 shows a flexible datacenter 200 powered by flare gas 800 in accordance with one or more embodiments of the present invention. Flare gas 800 is combustible gas produced as a product or by-product of petroleum refineries, chemical plants, natural gas processing plants, oil and gas drilling rigs, and oil and gas production facilities. Flare gas 800 is typically burned off through a flare stack (not shown) or vented into the air. In one or more embodiments of the present invention, flare gas 800 may be diverted 812 to a gas-powered generator that produces three-phase gas-generated AC voltage 822. This power may be considered behind-the-meter and is not subject to transmission and distribution costs. As such, one or more flexible datacenters 200 may be powered by three-phase gas-generated AC voltage. Specifically, the three-phase behind-the-meter AC voltage used to power flexible datacenter 200 may be three-phase gas-generated AC voltage 822. Accordingly, flexible datacenter 200 may reside behind-the-meter, avoid transmission and distribution costs, and may be dynamically powered when unutilized behind-the-meter power is available.

[0068] Figure 9 shows a method of dynamic power delivery to a flexible datacenter (200 of Figure 2) using unutilized behind-the-meter power 900 in accordance with one or more embodiments of the present invention. In step 910, the datacenter control system (220 of Figure 4), or the remote master control system (420 of Figure 4), may monitor unutilized behind-the-meter power availability. In certain embodiments, monitoring may include receiving information or an operational directive from the local station control system (410 of Figure 4) or the grid operator (440 of Figure 4) corresponding to unutilized behind-the-meter power availability.

[0069] In step 920, the datacenter control system (220 of Figure 4), or the remote master control system (420 of Figure 4), may determine when a datacenter ramp-up condition is met. In certain embodiments, the datacenter ramp-up condition may be met when there is sufficient behind-the-meter power availability and there is no operational directive from the local station to go offline or reduce power. In step 930, the datacenter control system (220 of Figure 4) may enable behind-the-meter

power delivery to one or more computing systems (100 of Figure 2). In step 940, once ramped-up, the datacenter control system (220 of Figure 4) or the remote master control system (420 of Figure 4) may direct one or more computing systems (100 of Figure 2) to perform predetermined computational operations. In certain embodiments, the predetermined computational operations may include the execution of one or more hashing functions.

[0070] While operational, the datacenter control system (220 of Figure 4), or the remote master control system (420 of Figure 4), may receive an operational directive to modulate power consumption. In certain embodiments, the operational directive may be a directive to reduce power consumption. In such embodiments, the datacenter control system (220 of Figure 4) or the remote master control system (420 of Figure 4) may dynamically reduce power delivery to one or more computing systems (100 of Figure 2) or dynamically reduce power consumption of one or more computing systems. In other embodiments, the operational directive may be a directive to provide a power factor correction factor. In such embodiments, the datacenter control system (220 of Figure 4) or the remote master control system (420 of Figure 4) may dynamically adjust power delivery to one or more computing systems (100 of Figure 2) to achieve a desired power factor correction factor. In still other embodiments, the operational directive may be a directive to go offline or power down. In such embodiments, the datacenter control system (220 of Figure 4) may disable power delivery to one or more computing systems (100 of Figure 2).

[0071] The datacenter control system (220 of Figure 4), or the remote master control system (420 of Figure 4), may determine when a datacenter ramp-down condition is met. In certain embodiments, the datacenter ramp-down condition may be met if there is insufficient or anticipated to be insufficient behind-the-meter power availability or there is an operational directive from the local station to go offline or reduce power. The datacenter control system (220 of Figure 4) may disable behind-the-meter power delivery to one or more computing systems (100 of Figure 2). Once ramped-down, the datacenter control system (220 of Figure 4) remains powered and in communication with the remote master control system (420 of Figure 4) so that it may dynamically power the flexible datacenter (200 of Figure 2) when conditions change.

[0072] One of ordinary skill in the art will recognize that a datacenter control system (220 of Figure 4) may dynamically modulate power delivery' to one or more

computing systems (100 of Figure 2) of a flexible datacenter (200 of Figure 2) based on unutilized behind-the-meter power availability or an operational directive. The flexible datacenter (200 of Figure 2) may transition between a fully powered down state (while the datacenter control system remains powered), a fully powered up state, and various intermediate states in between. In addition, flexible datacenter (200 of Figure 2) may have a blackout state, where all power consumption, including that of the datacenter control system (220 of Figure 4) is halted. However, once the flexible datacenter (200 of Figure 2) enters the blackout state, it will have to be manually rebooted to restore power to datacenter control system (220 of Figure 4). Local station conditions or operational directives may cause flexible datacenter (200 of Figure 2) to ramp-up, reduce power consumption, change power factor, or ramp-down

[0073] Advantages of one or more embodiments of the present invention may include one or more of the following:

[0074] In one or more embodiments of the present invention, a method and system for dynamic power delivery to a flexible datacenter using unutilized energy sources provides a green solution to two prominent problems: the exponential increase in power required for growing blockchain operations and the unutilized and typically wasted energy generated from renewable energy sources.

[0075] In one or more embodiments of the present invention, a method and system for dynamic power delivery to a flexible datacenter using unutilized energy sources allows for the rapid deployment of mobile datacenters to local stations. The mobile datacenters may be deployed on site, near the source of power generation, and receive unutilized behind-the-meter power when it is available.

[0076] In one or more embodiments of the present invention, a method and system for dynamic power delivery to a flexible datacenter using unutilized energy sources allows for the power delivery to the datacenter to be modulated based on conditions or an operational directive received from the local station or the grid operator.

[0077] In one or more embodiments of the present invention, a method and system for dynamic power delivery to a flexible datacenter using unutilized energy sources may dynamically adjust power consumption by ramping-up, ramping-down, or adjusting the power consumption of one or more computing systems within the flexible datacenter.

WO 2019/139632

PCT/US2018/017950

[0078] In one or more embodiments of the present invention, a method and system for dynamic power delivery to a flexible datacenter using unutilized energy sources may be powered by unutilized behind-the-meter power that is free from transmission and distribution costs. As such, the flexible datacenter may perform computational operations, such as hashing function operations, with little to no energy cost.

[0079] In one or more embodiments of the present invention, a method and system for dynamic power delivery to a flexible datacenter using unutilized energy sources provides a number of benefits to the hosting local station. The local station may use the flexible datacenter to adjust a load, provide a power factor correction, to offload power, or operate in a manner that invokes a production tax credit.

[0080] While the present invention has been described with respect to the above-noted embodiments, those skilled in the art, having the benefit of this disclosure, will recognize that other embodiments may be devised that are within the scope of the invention as disclosed herein. Accordingly, the scope of the invention should be limited only by the appended claims.

CLAIMS

What is claimed is:

1. A flexible datacenter comprising:
 - a mobile container;
 - a behind-the-meter power input system;
 - a power distribution system;
 - a datacenter control system;
 - a plurality of computing systems; and
 - a climate control system,wherein the datacenter control system modulates power delivery to the plurality of computing systems based on unutilized behind-the-meter power availability or an operational directive.
2. The flexible datacenter of claim 1, further comprising:
 - a remote master control system.
3. The flexible datacenter of claim 1, wherein the behind-the-meter power input system comprises an input configured to receive three-phase behind-the-meter AC voltage and a supervisory AC-to-AC step-down transformer configured to step down the three-phase behind-the-meter AC voltage to a single-phase supervisory nominal AC voltage or an input configured to receive single-phase supervisory nominal AC voltage from a local station or metered source.
4. The flexible datacenter of claim 3, wherein the behind-the-meter power input system provides the single-phase supervisory nominal AC voltage to the datacenter control system
5. The flexible datacenter of claim 1, wherein the behind-the-meter power input system comprises an input configured to receive three-phase behind-the-meter AC voltage and an operational AC-to-AC step-down transformer configured to controllably step down the three-phase behind-the-meter AC voltage to three-phase nominal AC voltage.

WO 2019/139632

PCT/US2018/017950

6. The flexible datacenter of claim 5, wherein the datacenter control system controllably enables or disables generation of the three-phase nominal AC voltage by the operational AC-to-AC step-down transformer.
7. The flexible datacenter of claim 1, wherein the behind-the-meter power input system provides three-phases of the three-phase nominal AC voltage to the power distribution system
8. The flexible datacenter of claim 7, wherein the power distribution system controllably provides a single phase of the three-phase nominal AC voltage to each computing system of the plurality of computing systems.
9. The flexible data center of claim 7, wherein the datacenter control system controllably selects which phase of the three-phase nominal AC voltage the power distribution system provides to each computing system of the plurality of computing systems
10. The flexible datacenter of claim 7, wherein the datacenter control system modulates a power factor correction factor by controllably adjusting which phase of the three-phase nominal AC voltage each computing system of the plurality of computing systems receive.
11. The flexible datacenter of claim 5, wherein the three-phase behind-the-meter AC voltage comprises a three-phase wind-generated AC voltage output by one or more wind turbines prior to an AC-to-AC step-up transformer that steps up the three-phase wind-generated AC voltage to a three-phase grid AC voltage.
12. The flexible datacenter of claim 5, wherein the three-phase behind-the-meter AC voltage comprises a three-phase solar-generated AC voltage output by a DC-to-AC inverter that inputs solar-generated DC voltage from one or more solar panels and prior to an AC-to-AC step-up transformer that steps up the three-phase solar-generated AC voltage to a three-phase grid AC voltage.

WO 2019/139632

PCT/US2018/017950

13. The flexible datacenter of claim 5, wherein the three-phase behind-the-meter AC voltage comprises a three-phase gas-generated AC voltage output by a generator that inputs combustible gas diverted from a flare or venting system
14. The flexible datacenter of claim 5, wherein the three-phase behind-the-meter AC voltage is a three-phase metered AC voltage.
15. The flexible datacenter of claim 1, wherein unutilized behind-the-meter power availability comprises one or more of excess local power generation at a local station level, excess local power generation that a grid cannot receive, local power generation subject to economic curtailment, local power generation subject to reliability curtailment, local power generation subject to power factor correction, low local power generation, start up local power generation situations, transient local power generation situations, or testing local power generation situations where there is an economic advantage to using local behind-the-meter power generation to power the flexible datacenter.
16. The flexible datacenter of claim 1, wherein an operational directive comprises one or more of a local station directive, a remote master control directive, or a grid directive.
17. The flexible datacenter of claim 1, wherein an operational directive comprises one or more of a dispatchability directive or a forecast directive.
18. The flexible datacenter of claim 1, wherein an operational directive comprises a workload directive based on actual behind-the-meter power availability or projected behind-the-meter power availability
19. The flexible datacenter of claim 2, wherein the remote master control system dynamically adjusts power delivery to the flexible datacenter based on a remote master control directive.
20. The flexible datacenter of claim 1, wherein the climate control system comprises a computer room air conditioner system, a computer room air handler system, an evaporative cooling system, a refrigerant cooling system, an immersive cooling system,

WO 2019/139632


PCT/US2018/017950

or any other suitable heat extraction system configured to operate the plurality of computing systems within their operational temperature range

21. The flexible datacenter of claim 1, wherein the mobile container comprises a storage container configured for placement on a ground surface.
22. The flexible datacenter of claim 1, wherein the mobile container comprises a storage trailer on wheels.
23. The flexible datacenter of claim 5, further comprising a battery system configured to convert the three-phase nominal AC voltage to DC nominal voltage and store power in a plurality of storage cells.
24. The flexible datacenter of claim 23, wherein the DC nominal voltage from the plurality of storage cells are converted via a DC-to-AC inverter to three-phase nominal AC voltage for flexible datacenter use
25. A method of dynamic power delivery to a flexible datacenter using unutilized behind-the-meter power comprising:
 - monitoring unutilized behind-the-meter power availability;
 - determining when a datacenter ramp-up condition is met;
 - enabling behind-the-meter power delivery to one or more computing systems when the datacenter ramp-up condition is met; and
 - directing the one or more computing systems to perform predetermined computational operations
26. The method of claim 25, further comprising;
 - determining when a datacenter ramp-down condition is met; and
 - disabling power delivery to one or more computing systems when the datacenter ramp-down condition is met.
27. The method of claim 25, further comprising:
 - receiving an operational directive to go offline; and
 - disabling power delivery to the one or more computing systems.

28. The method of claim 25, further comprising:
- receiving an operational directive to reduce power consumption, and
 - dynamically reducing power delivery to a subset of the one or more computing systems
29. The method of claim 25, further comprising:
- receiving an operational directive to reduce power consumption; and
 - dynamically reducing power consumption of the one or more computing systems.
30. The method of claim 25, further comprising:
- receiving an operational directive to provide power factor correction; and
 - dynamically adjusting power delivery to a subset of the one or more computing systems to achieve a desired power factor correction factor.
31. The method of claim 25, wherein unutilized behind-the-meter power availability comprises one or more of excess local power generation at a local station level, excess local power generation that a grid cannot receive, local power generation subject to economic curtailment, local power generation subject to reliability curtailment, local power generation subject to power factor correction, low local power generation, start up local power generation situations, transient local power generation situations, or testing local power generation situations where there is an economic advantage to using local behind-the-meter power generation.
32. The method of claim 25, wherein the datacenter ramp-up condition is met if there is sufficient behind-the-meter power availability and there is no operational directive from a local station to go offline.
33. The method of claim 25, wherein the datacenter ramp-down condition is met if there is insufficient behind-the-meter power availability or there is an operational directive from a local station to go offline.
34. The method of claim 25, wherein the predetermined computational operations comprise execution of one or more hashing functions.

1/8



2/8




FIG. 2

3/8




FIG. 3

4/8




FIG. 4

5/8





FIG. 5

6/8



7/8




FIG. 7

8/8

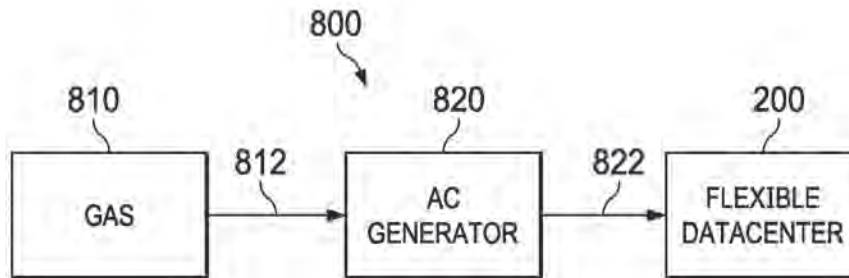


FIG. 8

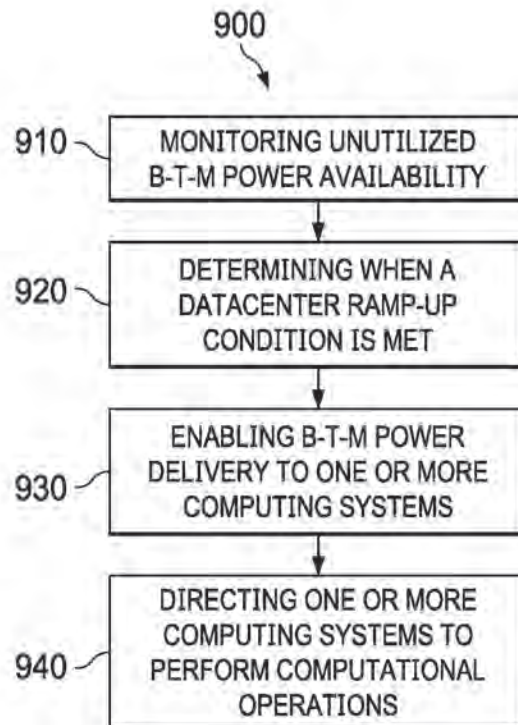


FIG. 9

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US2018/017950

A. CLASSIFICATION OF SUBJECT MATTER

IPC(8) - G06F 1/26; G06F 1/20; G06F 1/30; G06F 1/32 (2018.01)

CPC - G06F 1/26; H05K 7/1497; G06F 1/20; H05K 7/1485 (2018.05)

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

See Search History document

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

USPC - 713/300; 361/679.46; 361/679.02 (keyword delimited)

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

See Search History document

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2008/0094797 A1 (COGLITORE et al) 24 April 2008 (24.04.2008) entire document	1, 7, 15-18, 20-22
Y		2-6, 8-14, 19, 23, 24
Y	US 2010/0211810 A1 (ZACHO) 19 August 2010 (19.08.2010) entire document	2, 19
Y	US 2008/0030078 A1 (WHITTED et al) 07 February 2008 (07.02.2008) entire document	3-6, 8-14, 23, 24
Y	US 2013/0063991 A1 (XIAO et al) 14 March 2013 (14.03.2013) entire document	12
Y	US 2015/0155712 A1 (INERTECH IP LLC) 04 June 2015 (04.06.2015) entire document	24
A	US 2010/0328849 A1 (EWING et al) 30 December 2010 (30.12.2010) entire document	1-24
A	US 2013/0306276 A1 (DUCHESNEAU) 21 November 2013 (21.11.2013) entire document	1-24
A	US 2012/0300524 A1 (FORNAGE et al) 29 November 2012 (29.11.2012) entire document	1-24

☐ Further documents are listed in the continuation of Box C.☐ See patent family annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

14 May 2018

Date of mailing of the international search report

31 MAY 2018

Name and mailing address of the ISA/US

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents
P.O. Box 1450, Alexandria, VA 22313-1450

Facsimile No. 571-273-8300

Authorized officer

Blaine R. Copenheaver

PCT Helpdesk: 571-272-4300
PCT OSP: 571-272-7774

Form PCT/ISA/210 (second sheet) (January 2015)

LANCIUM00000091

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2018/017950

Box No. II Observations where certain claims were found unsearchable (Continuation of item 2 of first sheet)

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☐ Claims Nos.:
because they relate to subject matter not required to be searched by this Authority, namely:
2. ☐ Claims Nos.:
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:
3. ☐ Claims Nos.:
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box No. III Observations where unity of invention is lacking (Continuation of item 3 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:
See extra sheet(s).

1. ☐ As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.
2. ☐ As all searchable claims could be searched without effort justifying additional fees, this Authority did not invite payment of additional fees.
3. ☐ As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:
4. ☒ No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:
1-24

Remark on Protest

- ☐ The additional search fees were accompanied by the applicant's protest and, where applicable, the payment of a protest fee.
- ☐ The additional search fees were accompanied by the applicant's protest but the applicable protest fee was not paid within the time limit specified in the invitation.
- ☐ No protest accompanied the payment of additional search fees.

Form PCT/ISA/210, (continuation of first sheet (2)) (January 2015)

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US2018/017950

Continued from Box No. III Observations where unity of invention is lacking

This application contains the following inventions or groups of inventions which are not so linked as to form a single general inventive concept under PCT Rule 13.1. In order for all inventions to be examined, the appropriate additional examination fees must be paid.

Group I, claims 1-24, are drawn to a flexible datacenter comprising: a mobile container; and a behind-the-meter power input system.

Group II, claims 25-34, are drawn to a method of dynamic power delivery to a flexible datacenter using unutilized behind-the meter power comprising: monitoring unutilized behind-the-meter power availability; and determining when a datacenter ramp-up condition is met.

The inventions listed as Groups I-II do not relate to a single general inventive concept under PCT Rule 13.1 because, under PCT Rule 13.2, they lack the same or corresponding special technical features for the following reasons: the special technical feature of the Group I invention: a flexible datacenter comprising: a mobile container; a behind-the-meter power input system; a power distribution system; a datacenter control system; and a climate control system, wherein the datacenter control system modulates power delivery to the plurality of computing systems as claimed therein is not present in the invention of Group II. The special technical feature of the Group II invention: monitoring unutilized behind-the-meter power availability; determining when a datacenter ramp-up condition is met; enabling behind-the-meter power delivery to one or more computing systems when the datacenter ramp-up condition is met; and directing the one or more computing systems to perform predetermined computational operations as claimed therein is not present in the invention of Group I.

Groups I and II lack unity of invention because even though the inventions of these groups require the technical feature of a flexible datacenter comprising one or more computing systems; and power delivery to the flexible datacenter using unutilized behind-the meter power availability, this technical feature is not a special technical feature as it does not make a contribution over the prior art.

Specifically, US 2010/0328849 to Ewing teaches a datacenter comprising one or more computing systems; and power delivery to the datacenter using power availability (if a server demands 500 watts and the PUE for the datacenter is 3.0, then the power from the utility grid needed to deliver 500 watts to the server is 1500 watts. The DCIE, in comparison, may provide a different aspect of this information, a DCIE value of 0.33 (equivalent to a PUE of 3.0) suggesting that the computing equipment consumes 33% of the power in the data center, para.0060). Further, US 2013/0306276 to Duchesneau teaches a flexible datacenter (supercomputing datacenter, Para. [0315. Maximum flexibility, para. 1500]); and using unutilized behind-the meter power availability (PERKS is hybrid energy system combining UPS with a peak-shaving system that directly captures excess or low-cost energy from a multiplicity of sources (when it is cheapest or most readily available) and stores it for later reuse, such as during peak periods, Para. [1456]. Note that excess energy is hereby "unutilized behind the meter power" as defined by applicant in Para [0041] of the specification of the present application).

Since none of the special technical features of the Group I or II inventions are found in more than one of the inventions, unity of invention is lacking.

Exhibit 13



US011016456B2

(12) **United States Patent**
Henson et al.

(10) **Patent No.:** **US 11,016,456 B2**

(45) **Date of Patent:** ***May 25, 2021**

(54) **METHOD AND SYSTEM FOR DYNAMIC POWER DELIVERY TO A FLEXIBLE DATACENTER USING UNUTILIZED ENERGY SOURCES**

(71) Applicant: **Lancium LLC**, The Woodlands, TX (US)

(72) Inventors: **David Henson**, The Woodlands, TX (US); **Michael McNamara**, The Woodlands, TX (US); **Raymond Cline**, The Woodlands, TX (US)

(73) Assignee: **LANCIUM LLC**, Houston, TX (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

This patent is subject to a terminal disclaimer.

(21) Appl. No.: **16/482,495**

(22) PCT Filed: **Feb. 13, 2018**

(86) PCT No.: **PCT/US2018/017950**

§ 371 (c)(1),

(2) Date: **Jul. 31, 2019**

(87) PCT Pub. No.: **WO2019/139632**

PCT Pub. Date: **Jul. 18, 2019**

(65) **Prior Publication Data**

US 2020/0379537 A1 Dec. 3, 2020

Related U.S. Application Data

(60) Provisional application No. 62/616,348, filed on Jan. 11, 2018.

(51) **Int. Cl.**

G05B 19/042 (2006.01)

G06F 1/3203 (2019.01)

(Continued)

(52) **U.S. Cl.**

CPC **G05B 19/042** (2013.01); **A01G 9/18** (2013.01); **A01G 9/243** (2013.01); **A01G 9/246** (2013.01);

(Continued)

(58) **Field of Classification Search**

CPC **G05B 19/042**; **G05B 2219/2639**; **H02J 13/00002**; **H02J 3/381**; **H02J 13/00**;

(Continued)

(56) **References Cited**

U.S. PATENT DOCUMENTS

6,288,456 B1 9/2001 Cratty
6,633,823 B2 10/2003 Bartone et al.

(Continued)

FOREIGN PATENT DOCUMENTS

CN 103163904 A 6/2013
KR 20090012523 A 2/2009

(Continued)

OTHER PUBLICATIONS

Advisory Action dated Nov. 13, 2020 for U.S. Appl. No. 16/529,360, filed Aug. 1, 2019, 182 pages.

(Continued)


Primary Examiner — Michael J Brown

(74) *Attorney, Agent, or Firm* — Barnes & Thornburg LLP

(57) **ABSTRACT**

A flexible datacenter includes a mobile container, a behind-the-meter power input system, a power distribution system, a datacenter control system, a plurality of computing systems, and a climate control system. The datacenter control system modulates power delivery to the plurality of computing systems based on unutilized behind-the-meter power availability or an operational directive. A method of dynamic power delivery to a flexible datacenter using unutilized behind-the-meter power includes monitoring unutilized

(Continued)



US 11,016,456 B2

Page 2

lized behind-the-meter power availability, determining when a datacenter ramp-up condition is met, enabling behind-the-meter power delivery to one or more computing systems when the datacenter ramp-up condition is met, and directing the one or more computing systems to perform predetermined computational operations.

30 Claims, 8 Drawing Sheets

(51) Int. Cl.

H02J 3/38 (2006.01)
H02J 13/00 (2006.01)
H02S 40/32 (2014.01)
A01G 9/24 (2006.01)
A01G 9/18 (2006.01)
A01G 9/26 (2006.01)
G06F 9/48 (2006.01)

(52) U.S. Cl.

CPC *A01G 9/247* (2013.01); *A01G 9/249* (2019.05); *A01G 9/26* (2013.01); *G06F 1/3203* (2013.01); *G06F 9/4856* (2013.01); *H02J 3/381* (2013.01); *H02J 3/383* (2013.01); *H02J 3/386* (2013.01); *H02J 13/00* (2013.01); *H02J 13/00002* (2020.01); *H02J 13/0017* (2013.01); *H02S 40/32* (2014.12); *G05B 2219/2639* (2013.01); *H02J 2300/24* (2020.01); *H02J 2300/28* (2020.01)

(58) Field of Classification Search

CPC *H02J 3/383*; *H02J 3/386*; *H02J 13/0017*; *H02J 2300/28*; *H02J 2300/24*; *G06F 1/3203*; *G06F 9/4856*; *H02S 40/32*; *H02S 10/00*; *A01G 9/249*; *A01G 9/18*; *A01G 9/243*; *A01G 9/246*; *A01G 9/247*; *A01G 9/26*; *A01G 9/24*; *A01G 31/06*; *Y02E 10/76*; *Y02E 10/56*; *Y02P 60/21*; *Y02P 60/12*; *Y02A 40/25*
 USPC 713/300, 310, 320, 330
 See application file for complete search history.

(56)

References Cited

U.S. PATENT DOCUMENTS

7,143,300 B2 11/2006 Potter et al.
 7,647,516 B2 1/2010 Ranganathan et al.
 7,702,931 B2 4/2010 Goodrum et al.
 7,779,276 B2 8/2010 Bolan et al.
 7,861,102 B1 12/2010 Ranganathan et al.
 7,921,315 B2 4/2011 Langgood et al.
 7,970,561 B2 6/2011 Pfeiffer
 8,001,403 B2 8/2011 Hamilton et al.
 8,006,108 B2 8/2011 Brey et al.
 8,214,843 B2 7/2012 Boss et al.
 8,374,928 B2 2/2013 Gopisetty et al.
 8,447,993 B2 5/2013 Greene et al.
 8,571,820 B2 10/2013 Pfeiffer
 8,627,123 B2 1/2014 Jain et al.
 8,639,392 B2 1/2014 Chassin
 8,700,929 B1 4/2014 Weber et al.
 8,789,061 B2 7/2014 Pavel et al.
 8,799,690 B2 8/2014 Dawson et al.
 9,003,211 B2 4/2015 Pfeiffer
 9,003,216 B2 4/2015 Sankar et al.
 9,026,814 B2 5/2015 Aasheim et al.
 9,207,993 B2 12/2015 Jain
 9,218,035 B2 12/2015 Li et al.
 9,552,234 B2 1/2017 Boldyrev et al.
 9,645,596 B1 5/2017 Lee et al.

9,994,118 B2 6/2018 Williams et al.
 10,367,353 B1 7/2019 McNamara
 10,367,535 B2 7/2019 Corse et al.
 10,444,818 B1 10/2019 McNamara et al.
 10,452,127 B1 * 10/2019 McNamara G06F 1/3206
 10,497,072 B2 12/2019 Hooshmand et al.
 10,608,433 B1 3/2020 McNamara et al.
 10,618,427 B1 4/2020 McNamara et al.
 10,637,353 B2 4/2020 Ohyama et al.
 2002/0072868 A1 6/2002 Bartone et al.
 2002/0158749 A1 10/2002 Ikeda et al.
 2003/0074464 A1 4/2003 Bohrer et al.
 2004/0117330 A1 6/2004 Ehlers et al.
 2005/0203761 A1 9/2005 Barr et al.
 2006/0161765 A1 7/2006 Cromer et al.
 2008/0030078 A1 2/2008 Whitted et al.
 2008/0094797 A1 4/2008 Coglitore et al.
 2009/0055665 A1 2/2009 Maglione et al.
 2009/0070611 A1 3/2009 Bower, III et al.
 2009/0078401 A1 3/2009 Cichanowicz
 2009/0089595 A1 4/2009 Brey et al.
 2010/0211810 A1 8/2010 Zacho
 2010/0235004 A1 9/2010 Thind
 2010/0280675 A1 11/2010 Tate, Jr. et al.
 2010/0328849 A1 12/2010 Ewing et al.
 2011/0072289 A1 3/2011 Kato
 2011/0238342 A1 9/2011 Pfeiffer
 2012/0000121 A1 1/2012 Swann
 2012/0072745 A1 3/2012 Ahluwalia et al.
 2012/0300524 A1 11/2012 Fornage et al.
 2012/0306271 A1 12/2012 Kuriyama
 2012/0324259 A1 12/2012 Aasheim et al.
 2013/0006401 A1 1/2013 Shan
 2013/0063991 A1 3/2013 Xiao et al.
 2013/0086404 A1 4/2013 Sankar et al.
 2013/0187464 A1 7/2013 Smith et al.
 2013/0227139 A1 8/2013 Suffling
 2013/0306276 A1 11/2013 Duchesneau
 2014/0070756 A1 3/2014 Kearns et al.
 2014/0137468 A1 5/2014 Ching
 2014/0180886 A1 6/2014 Forbes, Jr.
 2014/0379156 A1 12/2014 Kamel et al.
 2015/0012113 A1 1/2015 Celebi
 2015/0121113 A1 4/2015 Ramamurthy et al.
 2015/0155712 A1 6/2015 Mondal
 2015/0212122 A1 7/2015 Sobotka et al.
 2015/0229227 A1 8/2015 Aeloiza et al.
 2015/0277410 A1 10/2015 Gupta et al.
 2015/0278968 A1 10/2015 Steven et al.
 2015/0288183 A1 10/2015 Villanueva, Jr. et al.
 2015/0372538 A1 12/2015 Siegler et al.
 2016/0006066 A1 1/2016 Robertson
 2016/0011617 A1 1/2016 Liu et al.
 2016/0043552 A1 2/2016 Villanueva, Jr. et al.
 2016/0126783 A1 5/2016 Cheng et al.
 2016/0170469 A1 6/2016 Sehgal et al.
 2016/0172900 A1 6/2016 Welch, Jr.
 2016/0187906 A1 6/2016 Bodas et al.
 2016/0198656 A1 7/2016 McNamara et al.
 2016/0212954 A1 7/2016 Argento
 2016/0324077 A1 11/2016 Frantzen et al.
 2017/0023969 A1 1/2017 Shows et al.
 2017/0104336 A1 4/2017 Elbsat et al.
 2017/0261949 A1 9/2017 Hoffmann et al.
 2017/0373500 A1 12/2017 Shafi et al.
 2018/0026478 A1 1/2018 Peloso
 2018/0144414 A1 5/2018 Lee et al.
 2018/0202825 A1 7/2018 You et al.
 2018/0240112 A1 8/2018 Castinado et al.
 2018/0366978 A1 12/2018 Mafan et al.
 2018/0367320 A1 12/2018 Montalvo
 2019/0052094 A1 2/2019 Pmsvsv et al.
 2019/0168630 A1 6/2019 Mrlik et al.
 2019/0258307 A1 8/2019 Shaikh et al.
 2019/0280521 A1 9/2019 Lundstrom et al.
 2019/0318327 A1 10/2019 Sowell et al.

(56)

References Cited

U.S. PATENT DOCUMENTS

2019/0324820 A1 10/2019 Krishnan et al.
2020/0073466 A1 3/2020 Walsh

FOREIGN PATENT DOCUMENTS

WO 2015199629 A1 12/2015
WO WO/2019139633 A1 7/2019

OTHER PUBLICATIONS

Advisory Action dated Oct. 22, 2020 for U.S. Appl. No. 16/528,348, filed Jul. 31, 2019, 3 pages.
Bakar.N et al., "Microgrid and Load Shedding Scheme During Islanded Mode: a Review," Elsevier, May 26, 2020, vol. 71, pp. 161-169. <https://www.sciencedirect.com/science/article/pii/S1364032116311030>.
Choi.Y et al., "Optimal Load Shedding for Maximizing Satisfaction in an Islanded Microgrid," Energies, 2017, vol. 10, pp. 45. doi: 10.3390/en10010045.
Final Office Action dated Jul. 23, 2020 on for U.S. Appl. No. 16/132,062, filed Sep. 14, 2018, 26 pages.
Final Office Action dated May 19, 2020 for U.S. Appl. No. 16/809,111, filed Mar. 4, 2020, 36 pages.
Final Office Action dated Jun. 3, 2020 for U.S. Appl. No. 16/528,348, filed Jul. 31, 2019, 33 pages.
Final Office Action dated May 28, 2020 for U.S. Appl. No. 16/132,098, filed Sep. 14, 2018, 24 pages.
Final Office Action dated Jul. 29, 2020 for U.S. Appl. No. 16/132,092, filed Sep. 14, 2018, 5 pages.
Gao.H et al., "Dynamic Load Shedding for an Islanded Microgrid With Limited Generation Resources," IET Generation, Transmission & Distribution, Sep. 2016, vol. 10(12), pp. 2953-2961. doi: 10.1049/iet-gtd.2015.1452.
International Search Report and Written Opinion of PCT Application No. PCT/US2020/044536, dated Aug. 26, 2020, 24 pages.
International Search Report and Written Opinion of PCT Application No. PCT/US2020/044539, dated Aug. 26, 2020, 7 pages.
John, "Stem and CPower to Combine Behind-the-Meter Batteries and Demand Response," Energy Storage, Aug. 8, 2017, 1 page.
Lim.Y et al., "Distributed Load-shedding System for Agent-based Autonomous Microgrid Operations," Energies, 2014, vol. 7(1), pp. 385-401. doi: 10.3390/en7010385.
Liu.W et al., "Improved Average Consensus Algorithm Based Distributed Cost Optimization for Loading Shedding of Autonomous Microgrids," International Journal of Electrical Power & Energy Systems, Dec. 2015, vol. 73, pp. 89-96. doi: 10.1016/j.ijepes.2015.04.006.
Mousavizadeh.S et al., "A Linear Two-stage Method for Resiliency Analysis in Distribution Systems Considering Renewable Energy and Demand Response Resources," Elsevier, 2017, pp. 443-460. doi: 10.1016/j.apenergy.2017.11.067.
Non-Final Office Action dated Jun. 12, 2020 for U.S. Appl. No. 16/803,109, filed Dec. 27, 2020, 31 pages.
Non-Final Office Action dated May 14, 2020 on for U.S. Appl. No. 16/834,987, filed Mar. 30, 2020, 30 pages.
Notice of Allowance dated Oct. 13, 2020 on for U.S. Appl. No. 16/132,098, filed Sep. 14, 2018, 5 pages.
Notice of Allowance dated Jun. 12, 2020 on for U.S. Appl. No. 16/834,987, filed Mar. 30, 2020, 9 pages.
Notice of Allowance dated Oct. 13, 2021 on for U.S. Appl. No. 16/175,246, filed Oct. 30, 2018, 5 pages.
Notice of Allowance dated Sep. 17, 2020 on for U.S. Appl. No. 16/175,246, filed Oct. 30, 2018, 5 pages.
Notice of Allowance dated Nov. 19, 2020 on for U.S. Appl. No. 16/132,062, filed Sep. 14, 2018, 7 pages.
Notice of Allowance dated Jan. 27, 2021 on for U.S. Appl. No. 16/132,092, filed Sep. 14, 2018, 8 pages.

Notice of Allowance dated Jul. 29, 2020 on for U.S. Appl. No. 16/132,011, filed Sep. 14, 2018, 5 pages.
Notice of Allowance dated Oct. 29, 2020 on for U.S. Appl. No. 16/132,092, filed Sep. 14, 2018, 8 pages.
Notice of Allowance dated Jan. 25, 2021 on for U.S. Appl. No. 16/132,098, filed Sep. 14, 2018, 5 pages.
Pashajavid.E et al., "A Multimode Supervisory Control Scheme for Coupling Remote Droop-Regulated Microgrids," IEEE Transactions on Smart Grid, May 26, 2020, vol. 9(5), pp. 5381-5392. <https://ieeexplore.ieee.org/abstract/document/7888570/>.
Pashajavid.E et al., "Frequency Support for Remote Microgrid Systems With Intermittent Distributed Energy Resources—A Two-level Hierarchical Strategy," IEEE Systems Journal, May 26, 2020, vol. 12(3), pp. 2760-2771. <https://ieeexplore.ieee.org/abstract/document/7862156/>.
Rudez.U and Mihalic.R, "Predictive Underfrequency Load Shedding Scheme for Islanded Power Systems With Renewable Generation," Electric Power Systems Research, May 2015, vol. 126, pp. 21-28. doi: 10.1016/j.epsr.2015.04.017.
Wilson, "Utility-Scale Deployment Project of Behind-the-Meter Energy Storage for Use in Ancillary Services, Energy Resiliency, Grid Infrastructure Investment Deferral, and Demand-Response Integration," Portland State University Thesis, 154 pages (Jun. 9, 2016).
Xu.Q et al., "Distributed Load Shedding for Microgrid With Compensation Support via Wireless Network," IET Generation, Transmission & Distribution, May 2018, vol. 12(9), pp. 2006-2018. doi: 10.1049/iet-gtd.2017.1029.
Zhou.Q et al., "Two-Stage Load Shedding for Secondary Control in Hierarchical Operation of Islanded Microgrids," IEEE Transactions on Smart Grid, May 2019, vol. 10(3), pp. 3103-3111. doi: 10.1109/TSG.2018.2817738.
Final Office Action dated Apr. 17, 2020 for U.S. Appl. No. 16/529,402, filed Aug. 1, 2019, 59 pages.
Hayes, Adam S., "A Cost of Production Model for Bitcoin," Department of Economics, The New School for Social Research, Mar. 2015, 5 pages.
Non-Final Office Action dated Feb. 20, 2020 for U.S. Appl. No. 16/702,894, filed Dec. 4, 2019, 30 pages.
Non-Final Office Action dated Apr. 2, 2020 on for U.S. Appl. No. 16/132,011, filed Sep. 14, 2018, 5 pages.
Non-Final Office Action dated Mar. 30, 2020 for U.S. Appl. No. 16/132,092, filed Sep. 14, 2018, 46 pages.
Notice of Allowance dated Mar. 2, 2020, for U.S. Appl. No. 16/596,190, filed Oct. 8, 2019, 15 pages.
Notice of Allowance dated Apr. 6, 2020, for U.S. Appl. No. 16/175,246, filed Oct. 30, 2018, 12 pages.
Final Office Action dated Oct. 1, 2019 for U.S. Appl. No. 16/175,246, filed Oct. 30, 2018, 18 pages.
Notice of Allowance dated Aug. 15, 2019, for U.S. Appl. No. 16/175,146, filed Oct. 30, 2018, 17 pages.
Wilson, J., "A Utility-Scale Deployment Project of Behind-the-Meter Energy Storage for Use in Ancillary Services, Energy Resiliency, Grid Infrastructure Investment Deferral, and Demand-Response Integration," Portland State University, 2016.
International Search Report and Written Opinion of PCT Application No. PCT/US20/57686, dated Feb. 22, 2021, 67 pages.
ISO, "Distributed Energy Resources Roadmap for New York's Wholesale Electricity Markets," Distributed Energy Resource Roadmap, Jan. 2017, pp. 1-39. [retrieved on Dec. 15, 2020]. Retrieved from the Internet: <url="https://www.nyiso.com/documents/20142/1391862/Distributed_Energy_Resources_Roadmap.pdf/ec0b3b64-4de2-73e0-ffef-49a4b8b1">https://www.nyiso.com/documents/20142/1391862/Distributed_Energy_Resources_Roadmap.pdf/ec0b3b64-4de2-73e0-ffef-49a4b8b1 b3ca.</url>.
Non-Final Office Action dated Feb. 4, 2021 on for U.S. Appl. No. 16/284,610, filed Feb. 25, 2019, 9 pages.
Non-Final Office Action dated Feb. 12, 2021 on for U.S. Appl. No. 16/528,348, filed Jul. 31, 2019, 54 pages.
Notice of Allowance dated Feb. 8, 2021 on for U.S. Appl. No. 16/132,062, filed Sep. 14, 2018, 21 pages.
Notice of Allowance dated Feb. 8, 2021 on for U.S. Appl. No. 16/803,109, filed Feb. 27, 2020, 29 pages.

US 11,016,456 B2

Page 4

(56)

References Cited

OTHER PUBLICATIONS

Notice of Allowance dated Feb. 8, 2021 on for U.S. Appl. No. 16/834,987, filed Mar. 30, 2020, 180 pages.
Notice of Allowance dated Jan. 25, 2021 on for U.S. Appl. No. 16/702,894, filed Dec. 4, 2019, 24 pages.
Co-pending U.S. Appl. No. 16/132,092, filed Sep. 14, 2018.
Co-pending U.S. Appl. No. 16/175,246, filed Oct. 30, 2018.
Notice of Allowance dated Jan. 27, 2020, for U.S. Appl. No. 16/702,931, filed Dec. 4, 2019, 23 pages.
Ghamkhari et al., "Optimal Integration of Renewable Energy Resources in Data Centers with Behind-the-Meter Renewable Generator," Department of Electrical and Computer Engineering Texas Tech University, 2012, pp. 3340-3444.
Non-Final Office Action dated Dec. 5, 2019 for U.S. Appl. No. 16/529,360, filed Aug. 1, 2019, 72 pages.
Non-Final Office Action dated Dec. 10, 2019 for U.S. Appl. No. 16/596,190, filed Oct. 8, 2019, 72 pages.
Non-Final Office Action dated Nov. 14, 2019 for U.S. Appl. No. 16/132,098, filed Sep. 14, 2018, 25 pages.
Non-Final Office Action dated Nov. 21, 2019 for U.S. Appl. No. 16/529,402, filed Aug. 1, 2019, 57 pages.

Non-Final Office Action dated Dec. 11, 2019 on for U.S. Appl. No. 16/132,062, filed Sep. 14, 2018, 17 pages.
Non-Final Office Action dated Dec. 10, 2019 for U.S. Appl. No. 16/528,348, filed Oct. 8, 2019, 33 pages.
Rahimi, Farrokh, "Using a Transactive Energy Framework," IEEE Electrification Magazine, Dec. 2016, pp. 23-29.
Bird L., et al., "Wind and Solar Energy Curtailment: Experience and Practices in the United States," National Renewable Energy Lab (NREL), Technical Report NREL/TP-6A20-60983, Mar. 2014, 58 pages.
EPEX Spot, "How They Occur, What They Mean," https://www.epexspot.com/en/company-info/basics_of_the_power_market/negative_prices, 2018, 2 pages.
International Search Report and Written Opinion of PCT Application No. PCT/US2018/017955, dated Apr. 30, 2018, 22 pages.
International Search Report and Written Opinion of PCT Application No. PCT/US2018/017950, dated May 31, 2018, 15 pages.
Notice of Allowance dated Apr. 2, 2019, for U.S. Appl. No. 16/175,335, filed Oct. 30, 2018, 12 pages.
Notice of Allowance dated Jul. 29, 2019, for U.S. Appl. No. 16/245,532, filed Jan. 11, 2019, 13 pages.
Soluna., "Powering the Block Chain," Aug. 2018, version 1.1, 29 pages.


* cited by examiner

U.S. Patent

May 25, 2021

Sheet 1 of 8

US 11,016,456 B2







FIG. 2

U.S. Patent

May 25, 2021

Sheet 3 of 8

US 11,016,456 B2



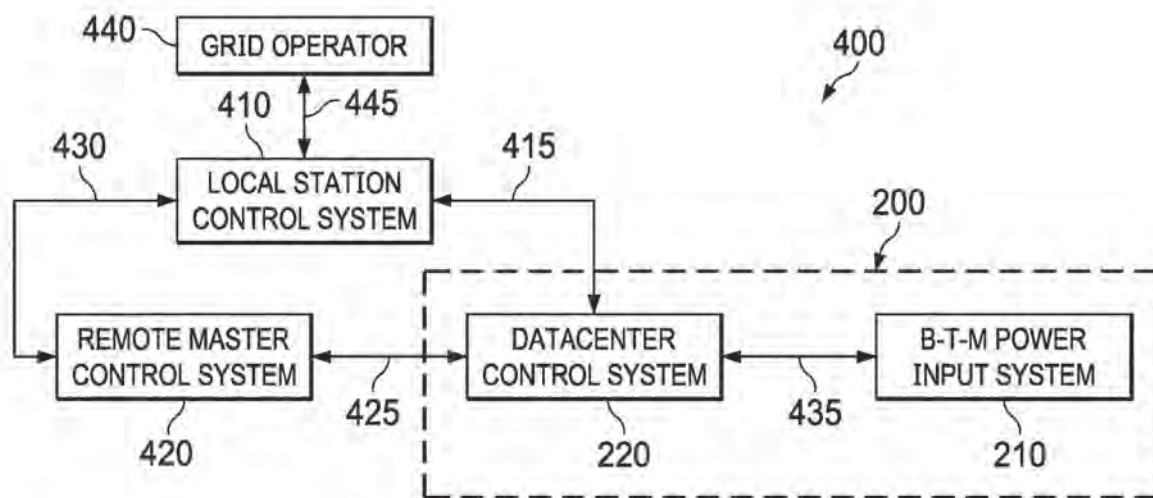


FIG. 4

U.S. Patent

May 25, 2021

Sheet 5 of 8

US 11,016,456 B2





FIG. 5

U.S. Patent

May 25, 2021

Sheet 6 of 8

US 11,016,456 B2






FIG. 7

U.S. Patent

May 25, 2021

Sheet 8 of 8

US 11,016,456 B2

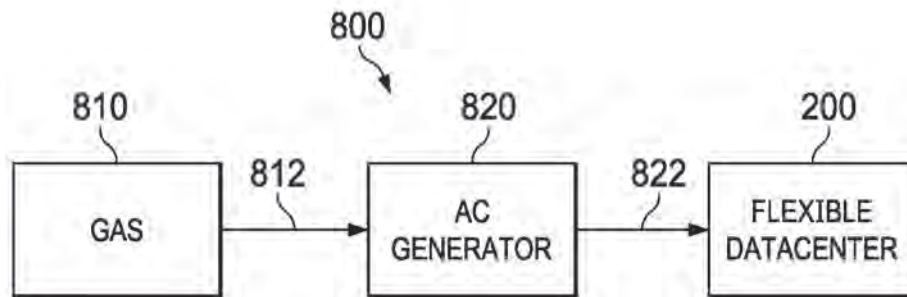


FIG. 8

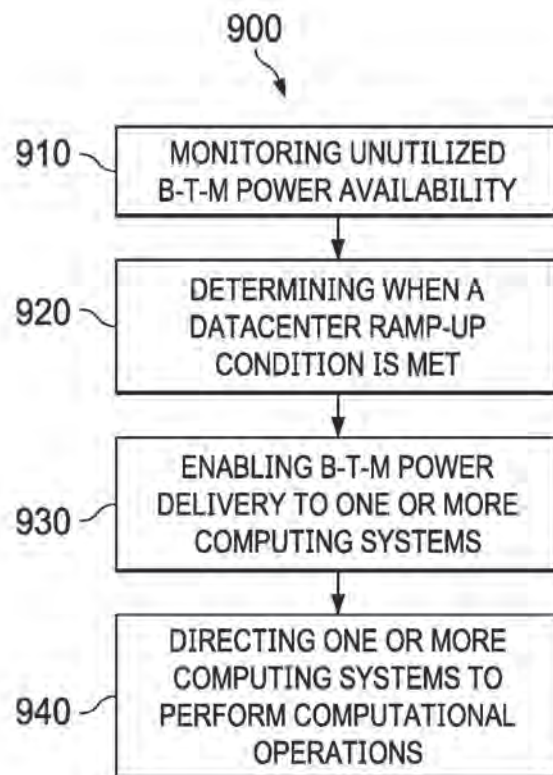


FIG. 9

US 11,016,456 B2

METHOD AND SYSTEM FOR DYNAMIC POWER DELIVERY TO A FLEXIBLE DATACENTER USING UNUTILIZED ENERGY SOURCES

CROSS-REFERENCE TO RELATED APPLICATIONS

This application is a U.S. national phase application under 35 USC 371 of International Patent Application no. PCT/US2018/017950, filed Feb. 13, 2018, which claims the benefit of U.S. Patent Application No. 62/616,348, filed Jan. 11, 2018.

BACKGROUND OF THE INVENTION

Blockchain technology was originally conceived of as an open and distributed system for securely conducting transactions with cryptographic currency. However, the foundational principle of blockchain technology is the ability to securely transact information of any type or kind between anonymous parties without intermediaries or a centralized trust authority. As such, blockchain technology finds application outside the realm of cryptocurrency and is widely considered one of the more robust and secure means of transacting information in the computer sciences.

In typical blockchain implementations, each participating party creates a digital identity, or wallet, which includes a pair of cryptographic keys used to transact information securely and anonymously with the blockchain. The blockchain may be thought of as a constantly growing database of all prior transaction information that is securely and coherently replicated across all nodes of a peer-to-peer blockchain network. The blockchain includes a sequence of blocks, where each block includes a bundle of transactions and other data including a hash of the prior block in the chain. As such, each block in the blockchain is mathematically related to the prior block, protecting the integrity of the blockchain from the most recently added block to the genesis block in the chain. Because anyone may participate in the curation of the blockchain, once a block is added, it becomes a permanent and immutable part of the blockchain. Thus, the blockchain stores transactions in a manner that prevents the transactions from being altered or otherwise corrupted, unless all subsequent blocks in the blockchain are also altered. The immutability of the blockchain makes the malicious alteration of a block exceptionally difficult, if not impossible, and at the very least makes it easy to detect and deter any such attempt before being accepted and replicated across the blockchain network.

Each transacting party of the blockchain uses a pair of cryptographic keys to anonymously transact information. The private key is a random number maintained in secrecy by the party holder that is used to derive a public key and sign information. The private key and the public key are mathematically related such that anyone holding the public key may verify that information signed with the private key originated from the holder of the private key. When an initiating party wishes to transact information, the information is signed with the initiating party's private key and broadcast to the blockchain network. A blockchain miner uses the initiating party's public key to verify that the initiating party initiated, or signed, the transaction. Once the initiating party's signature is validated, the transaction is validated, added to the next block in the blockchain, and replicated across all nodes.

The computational overhead of the blockchain is largely due to hashing functions used by blockchain miners to discover new blocks. While computationally intensive, the work performed by miners is critically important to the functionality of the blockchain. When an initiating party's transaction request has been lodged and the signatures validated, the transaction request is pooled in the blockchain network. Blockchain miners validate transactions and compete to discover a new block to be added to the blockchain. In order to add a newly discovered block to the blockchain, the blockchain miner must provide a cryptographic proof of the discovered block. To create the proof, the miner inputs the hash value of the prior block in the blockchain, the candidate block to be added, and a random number, commonly referred to as the nonce, to a hash function. The hash function takes input of any length and outputs an alphanumeric string of fixed length, commonly referred to as a hash, which uniquely identifies the input. However, the blockchain algorithm requires that the hash start with a certain number of leading zeros as determined by the current level of prescribed difficulty. The blockchain network modulates the level of difficulty for block discovery, by varying the number of leading zeros required in the calculated hash, based on the amount of computing power in the blockchain network.

As more computational capacity has come online, the hash rate has increased dramatically. In an effort to keep the block discovery time constant, the blockchain network modulates difficulty every 2016 blocks discovered. If the blockchain network hash rate is too high and the amount of time taken to discover a new block is less than 10 minutes, the difficulty is increased proportionally to increase the block discovery time to 10 minutes. Similarly, if the blockchain hash rate is too low and the amount of time taken to discover a new block is more than 10 minutes, the difficulty is increased proportionally to reduce the block discovery time to 10 minutes. Because there is no way to predict what hash value a given set of input data will generate, miners often have to execute the hash function a substantial number of times, each time inputting a new nonce, to generate a new hash value. When a miner is the first to obtain a hash value having the correct number of leading zeros, they broadcast the newly discovered block to the blockchain network and the blockchain is replicated across all nodes.

BRIEF SUMMARY OF THE INVENTION

According to one aspect of one or more embodiments of the present invention, a flexible datacenter includes a mobile container, a behind-the-meter power input system, a power distribution system, a datacenter control system, a plurality of computing systems, and a climate control system. The datacenter control system modulates power delivery to the plurality of computing systems based on unutilized behind-the-meter power availability or an operational directive.

According to one aspect of one or more embodiments of the present invention, a method of dynamic power delivery to a flexible datacenter using unutilized behind-the-meter power includes monitoring unutilized behind-the-meter power availability, determining when a datacenter ramp-up condition is met, enabling behind-the-meter power delivery to one or more computing systems when the datacenter ramp-up condition is met, and directing the one or more computing systems to perform predetermined computational operations.

US 11,016,456 B2

3

Other aspects of the present invention will be apparent from the following description and claims.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 shows a computing system in accordance with one or more embodiments of the present invention.

FIG. 2 shows a flexible datacenter in accordance with one or more embodiments of the present invention.

FIG. 3 shows a three-phase power distribution of a flexible datacenter in accordance with one or more embodiments of the present invention.

FIG. 4 shows a control distribution scheme of a flexible datacenter in accordance with one or more embodiments of the present invention.

FIG. 5 shows a control distribution scheme of a fleet of flexible datacenters in accordance with one or more embodiments of the present invention.

FIG. 6 shows a flexible datacenter powered by one or more wind turbines in accordance with one or more embodiments of the present invention.

FIG. 7 shows a flexible datacenter powered by one or more solar panels in accordance with one or more embodiments of the present invention.

FIG. 8 shows a flexible datacenter powered by flare gas in accordance with one or more embodiments of the present invention.

FIG. 9 shows a method of dynamic power delivery to a flexible datacenter using unutilized behind-the-meter power in accordance with one or more embodiments of the present invention.

DETAILED DESCRIPTION OF THE INVENTION

One or more embodiments of the present invention are described in detail with reference to the accompanying figures. For consistency, like elements in the various figures are denoted by like reference numerals. In the following detailed description of the present invention, specific details are set forth in order to provide a thorough understanding of the present invention. In other instances, well-known features to one having ordinary skill in the art are not described to avoid obscuring the description of the present invention.

Blockchain miners are typically compensated for their efforts through either a discovery fee or a fee paid by one or more of the transacting parties. Consequently, more and more computing resources are coming online to compete for these fees. As the number of computing resources increases, the blockchain network modulates the difficulty level, requiring hash values with more leading zeros. In essence, the increased difficulty means more hashing operations are required to find a valid hash. As such, there is an increasing number of computing resources executing an increasing number of hash functions that do not result in the discovery of a valid hash, yet still consume a substantial amount of power.

The intensive computational demand of blockchain applications makes the widespread adoption of blockchain technology inefficient and unsustainable from an energy and environmental perspective. In certain blockchain applications, with limited participation, roughly 5 quintillion 256-bit cryptographic hashes are created each and every second of every day. While it is difficult to determine how much energy is required for that computational task, it is estimated to be in excess of 500 megawatts, the vast majority of which is sourced from fossil fuels. The majority of blockchain

4

mining operations are currently being conducted in the People's Republic of China and powered by coal-fired energy. As blockchain technology proliferates, there is concern that the energy required to sustain such blockchain applications could exceed that of a developed country.

While future versions of blockchain technology may improve power consumption for various blockchain operations, including hashing functions, industry efforts have focused on the development of central processing units ("CPUs"), graphics processing units ("GPUs"), and application specific integrated circuits ("ASICs") that are specifically designed to perform blockchain operations in a more efficient manner. While such efforts are beneficial, the issue remains, the widespread adoption of blockchain technology will require substantially more power than is economically and environmentally feasible.

Accordingly, in one or more embodiments of the present invention, a method and system for dynamic power delivery to a flexible datacenter uses unutilized behind-the-meter power sources without transmission and distribution costs. The flexible datacenter may be configured to modulate power delivery to one or more computing systems based on the availability of unutilized behind-the-meter power or an operational directive. For example, the flexible datacenter may ramp-up to a fully online status, ramp-down to a fully offline status, or dynamically reduce power consumption, act as a load balancer, or adjust the power factor. Advantageously, the flexible datacenter may perform computational operations, such as blockchain hashing operations, with little to no energy costs, using clean and renewable energy that would otherwise be wasted.

FIG. 1 shows a computing system 100 in accordance with one or more embodiments of the present invention. Computing system 100 may include one or more central processing units (singular "CPU" or plural "CPUs") 105, host bridge 110, input/output ("IO") bridge 115, graphics processing units (singular "GPU" or plural "GPUs") 125, and/or application-specific integrated circuits (singular "ASIC" or plural "ASICs") (not shown) disposed on one or more printed circuit boards (not shown) that are configured to perform computational operations. Each of the one or more CPUs 105, GPUs 125, or ASICs (not shown) may be a single-core (not independently illustrated) device or a multi-core (not independently illustrated) device. Multi-core devices typically include a plurality of cores (not shown) disposed on the same physical die (not shown) or a plurality of cores (not shown) disposed on multiple die (not shown) that are collectively disposed within the same mechanical package (not shown).

CPU 105 may be a general purpose computational device typically configured to execute software instructions. CPU 105 may include an interface 108 to host bridge 110, an interface 118 to system memory 120, and an interface 123 to one or more IO devices, such as, for example, one or more GPUs 125. GPU 125 may serve as a specialized computational device typically configured to perform graphics functions related to frame buffer manipulation. However, one of ordinary skill in the art will recognize that GPU 125 may be used to perform non-graphics related functions that are computationally intensive. In certain embodiments, GPU 125 may interface 123 directly with CPU 125 (and interface 118 with system memory 120 through CPU 105). In other embodiments, GPU 125 may interface 121 with host bridge 110 (and interface 116 or 118 with system memory 120 through host bridge 110 or CPU 105 depending on the application or design). In still other embodiments, GPU 125 may interface 133 with IO bridge 115 (and interface 116 or

US 11,016,456 B2

5

118 with system memory 120 through host bridge 110 or CPU 105 depending on the application or design). The functionality of GPU 125 may be integrated, in whole or in part, with CPU 105.

Host bridge 110 may be an interface device configured to interface between the one or more computational devices and IO bridge 115 and, in some embodiments, system memory 120. Host bridge 110 may include an interface 108 to CPU 105, an interface 113 to IO bridge 115, for embodiments where CPU 105 does not include an interface 118 to system memory 120, an interface 116 to system memory 120, and for embodiments where CPU 105 does not include an integrated GPU 125 or an interface 123 to GPU 125, an interface 121 to GPU 125. The functionality of host bridge 110 may be integrated, in whole or in part, with CPU 105. IO bridge 115 may be an interface device configured to interface between the one or more computational devices and various IO devices (e.g., 140, 145) and IO expansion, or add-on, devices (not independently illustrated). IO bridge 115 may include an interface 113 to host bridge 110, one or more interfaces 133 to one or more IO expansion devices 135, an interface 138 to keyboard 140, an interface 143 to mouse 145, an interface 148 to one or more local storage devices 150, and an interface 153 to one or more network interface devices 155. The functionality of IO bridge 115 may be integrated, in whole or in part, with CPU 105 or host bridge 110. Each local storage device 150, if any, may be a solid-state memory device, a solid-state memory device array, a hard disk drive, a hard disk drive array, or any other non-transitory computer readable medium. Network interface device 155 may provide one or more network interfaces including any network protocol suitable to facilitate networked communications.

Computing system 100 may include one or more network-attached storage devices 160 in addition to, or instead of, one or more local storage devices 150. Each network-attached storage device 160, if any, may be a solid-state memory device, a solid-state memory device array, a hard disk drive, a hard disk drive array, or any other non-transitory computer readable medium. Network-attached storage device 160 may or may not be collocated with computing system 100 and may be accessible to computing system 100 via one or more network interfaces provided by one or more network interface devices 155.

One of ordinary skill in the art will recognize that computing system 100 may be a conventional computing system or an application-specific computing system. In certain embodiments, an application-specific computing system may include one or more ASICs (not shown) that are configured to perform one or more functions, such as hashing, in a more efficient manner. The one or more ASICs (not shown) may interface directly with CPU 105, host bridge 110, or GPU 125 or interface through IO bridge 115. Alternatively, in other embodiments, an application-specific computing system may be reduced to only those components necessary to perform a desired function in an effort to reduce one or more of chip count, printed circuit board footprint, thermal design power, and power consumption. The one or more ASICs (not shown) may be used instead of one or more of CPU 105, host bridge 110, IO bridge 115, or GPU 125. In such systems, the one or more ASICs may incorporate sufficient functionality to perform certain network and computational functions in a minimal footprint with substantially fewer component devices.

As such, one of ordinary skill in the art will recognize that CPU 105, host bridge 110, IO bridge 115, GPU 125, or ASIC (not shown) or a subset, superset, or combination of func-

6

tions or features thereof, may be integrated, distributed, or excluded, in whole or in part, based on an application, design, or form factor in accordance with one or more embodiments of the present invention. Thus, the description of computing system 100 is merely exemplary and not intended to limit the type, kind, or configuration of component devices that constitute a computing system 100 suitable for performing computing operations in accordance with one or more embodiments of the present invention.

One of ordinary skill in the art will recognize that computing system 100 may be a stand alone, laptop, desktop, server, blade, or rack mountable system and may vary based on an application or design.

FIG. 2 shows a flexible datacenter 200 in accordance with one or more embodiments of the present invention. Flexible datacenter 200 may include a mobile container 205, a behind-the-meter power input system 210, a power distribution system 215, a climate control system (e.g., 250, 260, 270, 280, and/or 290), a datacenter control system 220, and a plurality of computing systems 100 disposed in one or more racks 240. Datacenter control system 220 may be a computing system (e.g., 100 of FIG. 1) configured to dynamically modulate power delivery to one or more computing systems 100 disposed within flexible datacenter 200 based on unutilized behind-the-meter power availability or an operational directive from a local station control system (not shown), a remote master control system (not shown), or a grid operator (not shown).

In certain embodiments, mobile container 205 may be a storage trailer disposed on wheels and configured for rapid deployment. In other embodiments, mobile container 205 may be a storage container (not shown) configured for placement on the ground and potentially stacked in a vertical manner (not shown). In still other embodiments, mobile container 205 may be an inflatable container, a floating container, or any other type or kind of container suitable for housing a mobile datacenter 200.

Flexible datacenter 200 may be rapidly deployed on site near a source of unutilized behind-the-meter power generation. Behind-the-meter power input system 210 may be configured to input power to flexible datacenter 200. Behind-the-meter power input system 210 may include a first input (not independently illustrated) configured to receive three-phase behind-the-meter alternating current ("AC") voltage. In certain embodiments, behind-the-meter power input system 210 may include a supervisory AC-to-AC step-down transformer (not shown) configured to step down three-phase behind-the-meter AC voltage to single-phase supervisory nominal AC voltage or a second input (not independently illustrated) configured to receive single-phase supervisory nominal AC voltage from the local station (not shown) or a metered source (not shown). Behind-the-meter power input system 210 may provide single-phase supervisory nominal AC voltage to datacenter control system 220, which may remain powered at almost all times to control the operation of flexible datacenter 200. The first input (not independently illustrated) or a third input (not independently illustrated) of behind-the-meter power input system 210 may direct three-phase behind-the-meter AC voltage to an operational AC-to-AC step-down transformer (not shown) configured to controllably step down three-phase behind-the-meter AC voltage to three-phase nominal AC voltage. Datacenter control system 220 may controllably enable or disable generation or provision of three-phase nominal AC voltage by the operational AC-to-AC step-down transformer (not shown).

US 11,016,456 B2

7

Behind-the-meter power input system 210 may provide three phases of three-phase nominal AC voltage to power distribution system 215. Power distribution system 215 may controllably provide a single phase of three-phase nominal AC voltage to each computing system 100 or group 240 of computing systems 100 disposed within flexible datacenter 200. Datacenter control system 220 may controllably select which phase of three-phase nominal AC voltage that power distribution system 215 provides to each computing system 100 or group 240 of computing systems 100. In this way, datacenter control system 220 may modulate power delivery by either ramping-up flexible datacenter 200 to fully operational status, ramping-down flexible datacenter 200 to offline status (where only datacenter control system 220 remains powered), reducing power consumption by withdrawing power delivery from, or reducing power to, one or more computing systems 100 or groups 240 of computing systems 100, or modulating a power factor correction factor for the local station by controllably adjusting which phases of three-phase nominal AC voltage are used by one or more computing systems 100 or groups 240 of computing systems 100.

Flexible datacenter 200 may include a climate control system (e.g., 250, 260, 270, 280, 290) configured to maintain the plurality of computing systems 100 within their operational temperature range. In certain embodiments, the climate control system may include an air intake 250, an evaporative cooling system 270, a fan 280, and an air outtake 260. In other embodiments, the climate control system may include an air intake 250, an air conditioner or refrigerant cooling system 290, and an air outtake 260. In still other embodiments, the climate control system may include a computer room air conditioner system (not shown), a computer room air handler system (not shown), or an immersive cooling system (not shown). One of ordinary skill in the art will recognize that any suitable heat extraction system (not shown) configured to maintain the operation of the plurality of computing systems 100 within their operational temperature range may be used in accordance with one or more embodiments of the present invention.

Flexible datacenter 200 may include a battery system (not shown) configured to convert three-phase nominal AC voltage to nominal DC voltage and store power in a plurality of storage cells. The battery system (not shown) may include a DC-to-AC inverter configured to convert nominal DC voltage to three-phase nominal AC voltage for flexible datacenter 200 use. Alternatively, the battery system (not shown) may include a DC-to-AC inverter configured to convert nominal DC voltage to single-phase nominal AC voltage to power datacenter control system 220.

One of ordinary skill in the art will recognize that a voltage level of three-phase behind-the-meter AC voltage may vary based on an application or design and the type or kind of local power generation. As such, a type, kind, or configuration of the operational AC-to-AC step down transformer (not shown) may vary based on the application or design. In addition, the frequency and voltage level of three-phase nominal AC voltage, single-phase nominal AC voltage, and nominal DC voltage may vary based on the application or design in accordance with one or more embodiments of the present invention.

FIG. 3 shows a three-phase power distribution of a flexible datacenter 200 in accordance with one or more embodiments of the present invention. Flexible datacenter 200 may include a plurality of racks 240, each of which may include one or more computing systems 100 disposed therein. As discussed above, the behind-the-meter power

8

input system (210 of FIG. 2) may provide three phases of three-phase nominal AC voltage to the power distribution system (215 of FIG. 2). The power distribution system (215 of FIG. 2) may controllably provide a single phase of three-phase nominal AC voltage to each computing system 100 or group 240 of computing systems 100 disposed within flexible datacenter 200. For example, a flexible datacenter 200 may include eighteen racks 240, each of which may include eighteen computing systems 100. The power distribution system (215 of FIG. 2) may control which phase of three-phase nominal AC voltage is provided to one or more computing systems 100, a rack 240 of computing systems 100, or a group (e.g., 310, 320, or 330) of racks 240 of computing systems 100.

In the figure, for purposes of illustration only, eighteen racks 240 are divided into a first group of six racks 310, a second group of six racks 320, and a third group of six racks 330, where each rack contains eighteen computing systems 100. The power distribution system (215 of FIG. 2) may, for example, provide a first phase of three-phase nominal AC voltage to the first group of six racks 310, a second phase of three-phase nominal AC voltage to the second group of six racks 320, and a third phase of three-phase nominal AC voltage to the third group of six racks 330. If the flexible datacenter (200 of FIG. 2) receives an operational directive from the local station (not shown) to provide power factor correction, the datacenter control system (220 of FIG. 2) may direct the power distribution system (215 of FIG. 2) to adjust which phase or phases of three-phase nominal AC voltage are used to provide the power factor correction required by the local station (not shown) or grid operator (not shown). One of ordinary skill in the art will recognize that, in addition to the power distribution, the load may be varied by adjusting the number of computing systems 100 operatively powered. As such, the flexible datacenter (200 of FIG. 2) may be configured to act as a capacitive or inductive load to provide the appropriate reactance necessary to achieve the power factor correction required by the local station (not shown).

FIG. 4 shows a control distribution scheme of a flexible datacenter 200 in accordance with one or more embodiments of the present invention. Datacenter control system 220 may independently, or cooperatively with one or more of local station control system 410, remote master control system 420, and grid operator 440, modulate power delivery to flexible datacenter 200. Specifically, power delivery may be dynamically adjusted based on conditions or operational directives.

Local station control system 410 may be a computing system (e.g., 100 of FIG. 1) that is configured to control various aspects of the local station (not independently illustrated) that generates power and sometimes generates unutilized behind-the-meter power. Local station control system 410 may communicate with remote master control system 420 over a networked connection 430 and with datacenter control system 220 over a networked or hardwired connection 415. Remote master control system 420 may be a computing system (e.g., 100 of FIG. 1) that is located offsite, but connected via a network connection 425 to datacenter control system 220, that is configured to provide supervisory or override control of flexible datacenter 200 or a fleet (not shown) of flexible datacenters 200. Grid operator 440 may be a computing system (e.g., 100 of FIG. 1) that is configured to control various aspects of the grid (not independently illustrated) that receives power from the local station (not independently illustrated). Grid operator 440 may commu-

nicate with local station control system 440 over a networked or hardwired connection 445.

Datacenter control system 220 may monitor unutilized behind-the-meter power availability at the local station (not independently illustrated) and determine when a datacenter ramp-up condition is met. Unutilized behind-the-meter power availability may include one or more of excess local power generation, excess local power generation that the grid cannot accept, local power generation that is subject to economic curtailment, local power generation that is subject to reliability curtailment, local power generation that is subject to power factor correction, situations where local power generation is prohibitively low, start up situations, transient situations, or testing situations where there is an economic advantage to using locally generated behind-the-meter power generation, specifically power available at little to no cost and with no associated transmission or distribution costs.

The datacenter ramp-up condition may be met if there is sufficient behind-the-meter power availability and there is no operational directive from local station control system 410, remote master control system 420, or grid operator 440 to go offline or reduce power. As such, datacenter control system 220 may enable 435 behind-the-meter power input system 210 to provide three-phase nominal AC voltage to the power distribution system (215 of FIG. 2) to power the plurality of computing systems (100 of FIG. 2) or a subset thereof. Datacenter control system 220 may optionally direct one or more computing systems (100 of FIG. 2) to perform predetermined computational operations. For example, if the one or more computing systems (100 of FIG. 2) are configured to perform blockchain hashing operations, datacenter control system 220 may direct them to perform blockchain hashing operations for a specific blockchain application, such as, for example, Bitcoin, Litecoin, or Ethereum. Alternatively, one or more computing systems (100 of FIG. 2) may be configured to independently receive a computational directive from a network connection (not shown) to a peer-to-peer blockchain network (not shown) such as, for example, a network for a specific blockchain application, to perform predetermined computational operations.

Remote master control system 420 may specify to datacenter control system 220 what sufficient behind-the-meter power availability constitutes, or datacenter control system 220 may be programmed with a predetermined preference or criteria on which to make the determination independently. For example, in certain circumstances, sufficient behind-the-meter power availability may be less than that required to fully power the entire flexible datacenter 200. In such circumstances, datacenter control system 220 may provide power to only a subset of computing systems (100 of FIG. 2), or operate the plurality of computing systems (100 of FIG. 2) in a lower power mode, that is within the sufficient, but less than full, range of power that is available.

While flexible datacenter 200 is online and operational, a datacenter ramp-down condition may be met when there is insufficient, or anticipated to be insufficient, behind-the-meter power availability or there is an operational directive from local station control system 410, remote master control system 420, or grid operator 440. Datacenter control system 220 may monitor and determine when there is insufficient, or anticipated to be insufficient, behind-the-meter power availability. As noted above, sufficiency may be specified by remote master control system 420 or datacenter control system 220 may be programmed with a predetermined preference or criteria on which to make the determination

independently. An operational directive may be based on current dispatchability, forward looking forecasts for when unutilized behind-the-meter power is, or is expected to be, available, economic considerations, reliability considerations, operational considerations, or the discretion of the local station 410, remote master control 420, or grid operator 440. For example, local station control system 410, remote master control system 420, or grid operator 440 may issue an operational directive to flexible datacenter 200 to go offline and power down. When the datacenter ramp-down condition is met, datacenter control system 220 may disable power delivery to the plurality of computing systems (100 of FIG. 2). Datacenter control system 220 may disable 435 behind-the-meter power input system 210 from providing three-phase nominal AC voltage to the power distribution system (215 of FIG. 2) to power down the plurality of computing systems (100 of FIG. 2), while datacenter control system 220 remains powered and is capable of rebooting flexible datacenter 200 when unutilized behind-the-meter power becomes available again.

While flexible datacenter 200 is online and operational, changed conditions or an operational directive may cause datacenter control system 220 to modulate power consumption by flexible datacenter 200. Datacenter control system 220 may determine, or local station control system 410, remote master control system 420, or grid operator 440 may communicate, that a change in local conditions may result in less power generation, availability, or economic feasibility, than would be necessary to fully power flexible datacenter 200. In such situations, datacenter control system 220 may take steps to reduce or stop power consumption by flexible datacenter 200 (other than that required to maintain operation of datacenter control system 220). Alternatively, local station control system 410, remote master control system 420, or grid operator 440, may issue an operational directive to reduce power consumption for any reason, the cause of which may be unknown. In response, datacenter control system 220 may dynamically reduce or withdraw power delivery to one or more computing systems (100 of FIG. 2) to meet the dictate. Datacenter control system 220 may controllably provide three-phase nominal AC voltage to a smaller subset of computing systems (100 of FIG. 2) to reduce power consumption. Datacenter control system 220 may dynamically reduce the power consumption of one or more computing systems (100 of FIG. 2) by reducing their operating frequency or forcing them into a lower power mode through a network directive.

One of ordinary skill in the art will recognize that datacenter control system 220 may be configured to have a number of different configurations, such as a number or type or kind of computing systems (100 of FIG. 2) that may be powered, and in what operating mode, that correspond to a number of different ranges of sufficient and available unutilized behind-the-meter power availability. As such, datacenter control system 220 may modulate power delivery over a variety of ranges of sufficient and available unutilized behind-the-meter power availability.

FIG. 5 shows a control distribution of a fleet 500 of flexible datacenters 200 in accordance with one or more embodiments of the present invention. The control distribution of a flexible datacenter 200 shown and described with respect to FIG. 4 may be extended to a fleet 500 of flexible datacenters 200. For example, a first local station (not independently illustrated), such as, for example, a wind farm (not shown), may include a first plurality 510 of flexible datacenters 200a through 200d, which may be collocated or distributed across the local station (not shown). A second

US 11,016,456 B2

11

local station (not independently illustrated), such as, for example, another wind farm or a solar farm (not shown), may include a second plurality 520 of flexible datacenters 200e through 200h, which may be collocated or distributed across the local station (not shown). One of ordinary skill in the art will recognize that the number of flexible datacenters 200 deployed at a given station and the number of stations within the fleet may vary based on an application or design in accordance with one or more embodiments of the present invention.

Remote master control system 420 may provide supervisory control over fleet 500 of flexible datacenters 200 in a similar manner to that shown and described with respect to FIG. 4, with the added flexibility to make high level decisions with respect to fleet 500 that may be counterintuitive to a given station. Remote master control system 420 may make decisions regarding the issuance of operational directives to a given local station based on, for example, the status of each local station where flexible datacenters 200 are deployed, the workload distributed across fleet 500, and the expected computational demand required for the expected workload. In addition, remote master control system 420 may shift workloads from a first plurality 510 of flexible datacenters 200 to a second plurality 520 of flexible datacenters 200 for any reason, including, for example, a loss of unutilized behind-the-meter power availability at one local station and the availability of unutilized behind-the-meter power at another local station.

FIG. 6 shows a flexible datacenter 200 powered by one or more wind turbines 610 in accordance with one or more embodiments of the present invention. A wind farm 600 typically includes a plurality of wind turbines 610, each of which intermittently generates a wind-generated AC voltage. The wind-generated AC voltage may vary based on a type, kind, or configuration of farm 600, turbine 610, and incident wind speed. The wind-generated AC voltage is typically input into a turbine AC-to-AC step-up transformer (not shown) that is disposed within the nacelle (not independently illustrated) or at the base of the mast (not independently illustrated) of turbine 610. The turbine AC-to-AC step up transformer (not shown) outputs three-phase wind-generated AC voltage 620. Three-phase wind-generated AC voltage 620 produced by the plurality of wind turbines 610 is collected 625 and provided 630 to another AC-to-AC step-up transformer 640 that steps up three-phase wind-generated AC voltage 620 to three-phase grid AC voltage 650 suitable for delivery to grid 660. Three-phase grid AC voltage 650 may be stepped down with an AC-to-AC step-down transformer 670 configured to produce three-phase local station AC voltage 680 provided to local station 690. One of ordinary skill in the art will recognize that the actual voltage levels may vary based on the type, kind, or number of wind turbines 610, the configuration or design of wind farm 600, and grid 660 that it feeds into.

The output side of AC-to-AC step-up transformer 640 that connects to grid 660 may be metered and is typically subject to transmission and distribution costs. In contrast, power consumed on the input side of AC-to-AC step-up transformer 640 may be considered behind-the-meter and is typically not subject to transmission and distribution costs. As such, one or more flexible datacenters 200 may be powered by three-phase wind-generated AC voltage 620. Specifically, in wind farm 600 applications, the three-phase behind-the-meter AC voltage used to power flexible datacenter 200 may be three-phase wind-generated AC voltage 620. As such, flexible datacenter 200 may reside behind-

12

the-meter, avoid transmission and distribution costs, and may be dynamically powered when unutilized behind-the-meter power is available.

Unutilized behind-the-meter power availability may occur when there is excess local power generation. In high wind conditions, wind farm 600 may generate more power than, for example, AC-to-AC step-up transformer 640 is rated for. In such situations, wind farm 600 may have to take steps to protect its equipment from damage, which may include taking one or more turbines 610 offline or shunting their voltage to dummy loads or ground. Advantageously, one or more flexible datacenters 200 may be used to consume power on the input side of AC-to-AC step-up transformer 640, thereby allowing wind farm 600 to operate equipment within operating ranges while flexible datacenter 200 receives behind-the-meter power without transmission or distribution costs. The local station control system (not independently illustrated) of local station 690 may issue an operational directive to the one or more flexible datacenters 200 or to the remote master control system (420 of FIG. 4) to ramp-up to the desired power consumption level. When the operational directive requires the cooperative action of multiple flexible datacenters 200, the remote master control system (420 of FIG. 4) may determine how to power each individual flexible datacenter 200 in accordance with the operational directive or provide an override to each flexible datacenter 200.

Another example of unutilized behind-the-meter power availability is when grid 660 cannot, for whatever reason, take the power being produced by wind farm 600. In such situations, wind farm 600 may have to take one or more turbines 610 offline or shunt their voltage to dummy loads or ground. Advantageously, one or more flexible datacenters 200 may be used to consume power on the input side of AC-to-AC step-up transformer 640, thereby allowing wind farm 600 to either produce power to grid 660 at a lower level or shut down transformer 640 entirely while flexible datacenter 200 receives behind-the-meter power without transmission or distribution costs. The local station control system (not independently illustrated) of local station 690 or the grid operator (not independently illustrated) of grid 660 may issue an operational directive to the one or more flexible datacenters 200 or to the remote master control system (420 of FIG. 4) to ramp-up to the desired power consumption level. When the operational directive requires the cooperative action of multiple flexible datacenters 200, the remote master control system (420 of FIG. 4) may determine how to power each individual flexible datacenter 200 in accordance with the operational directive or provide an override to each flexible datacenter 200.

Another example of unutilized behind-the-meter power availability is when wind farm 600 is selling power to grid 660 at a negative price that is offset by a production tax credit. In certain circumstances, the value of the production tax credit may exceed the price wind farm 600 would have to pay to grid 660 to offload their generated power. Advantageously, one or more flexible datacenters 200 may be used to consume power behind-the-meter, thereby allowing wind farm 600 to produce and obtain the production tax credit, but sell less power to grid 660 at the negative price. The local station control system (not independently illustrated) of local station 690 may issue an operational directive to the one or more flexible datacenters 200 or to the remote master control system (420 of FIG. 4) to ramp-up to the desired power consumption level. When the operational directive requires the cooperative action of multiple flexible datacenter 200, the remote master control system (420 of FIG. 4)

US 11,016,456 B2

13

may determine how to power each individual flexible datacenter 200 in accordance with the operational directive or provide an override to each flexible datacenter 200.

Another example of unutilized behind-the-meter power availability is when wind farm 600 is selling power to grid 660 at a negative price because grid 660 is oversupplied or is instructed to stand down and stop producing altogether. The grid operator (not independently illustrated) may select certain power generation stations to go offline and stop producing power to grid 660. Advantageously, one or more flexible datacenters 200 may be used to consume power behind-the-meter, thereby allowing wind farm 600 to stop producing power to grid 660, but making productive use of the power generated behind-the-meter without transmission or distribution costs. The local station control system (not independently illustrated) of the local station 690 or the grid operator (not independently illustrated) of grid 660 may issue an operational directive to the one or more flexible datacenters 200 or to the remote master control system (420 of FIG. 4) to ramp-up to the desired power consumption level. When the operational directive requires the cooperative action of multiple flexible datacenters 200, the remote master control system (420 of FIG. 4) may determine how to power each individual flexible datacenter 200 in accordance with the operational directive or provide an override to each flexible datacenter 200.

Another example of unutilized behind-the-meter power availability is when wind farm 600 is producing power to grid 660 that is unstable, out of phase, or at the wrong frequency, or grid 660 is already unstable, out of phase, or at the wrong frequency for whatever reason. The grid operator (not independently illustrated) may select certain power generation stations to go offline and stop producing power to grid 660. Advantageously, one or more flexible datacenters 200 may be used to consume power behind-the-meter, thereby allowing wind farm 600 to stop producing power to grid 660, but make productive use of the power generated behind-the-meter without transmission or distribution costs. The local station control system (not independently illustrated) of local station 690 may issue an operational directive to the one or more flexible datacenters 200 or to the remote master control system (420 of FIG. 4) to ramp-up to the desired power consumption level. When the operational directive requires the cooperative action of multiple flexible datacenters 200, the remote master control system (420 of FIG. 4) may determine how to power each individual flexible datacenter 200 in accordance with the operational directive or provide an override to each flexible datacenter 200.

Further examples of unutilized behind-the-meter power availability is when wind farm 600 experiences low wind conditions that make it not economically feasible to power up certain components, such as, for example, the local station (not independently illustrated), but there may be sufficient behind-the-meter power availability to power one or more flexible datacenters 200. Similarly, unutilized behind-the-meter power availability may occur when wind farm 600 is starting up, or testing, one or more turbines 610. Turbines 610 are frequently offline for installation, maintenance, and service and must be tested prior to coming online as part of the array. One or more flexible datacenters 200 may be powered by one or more turbines 610 that are offline from farm 600. The above-noted examples of when unutilized behind-the-meter power is available are merely exemplary and are not intended to limit the scope of what one of ordinary skill in the art would recognize as unutilized behind-the-meter power availability. Unutilized behind-the-

14

meter power availability may occur anytime there is power available and accessible behind-the-meter that is not subject to transmission and distribution costs and there is an economic advantage to using it.

One of ordinary skill in the art will recognize that wind farm 600 and wind turbine 610 may vary based on an application or design in accordance with one or more embodiments of the present invention.

FIG. 7 shows a flexible datacenter 200 powered by one or more solar panels 710 in accordance with one or more embodiments of the present invention. A solar farm 700 typically includes a plurality of solar panels 710, each of which intermittently generates a solar-generated DC voltage 720. Solar-generated DC voltage 720 may vary based on a type, kind, or configuration of farm 700, panel 710, and incident sunlight. Solar-generated DC voltage 720 produced by the plurality of solar panels 710 is collected 725 and provided 730 to a DC-to-AC inverter that converts solar-generated DC voltage into three-phase solar-generated AC voltage 750. Three-phase solar-generated AC voltage 750 is provided to an AC-to-AC step-up transformer 760 that steps up three-phase solar-generated AC voltage to three-phase grid AC voltage 790. Three-phase grid AC voltage 790 may be stepped down with an AC-to-AC step-down transformer 785 configured to produce three-phase local station AC voltage 777 provided to local station 775. One of ordinary skill in the art will recognize that the actual voltage levels may vary based on the type, kind, or number of solar panels 710, the configuration or design of solar farm 700, and grid 790 that it feeds into.

The output side of AC-to-AC step-up transformer 760 that connects to grid 790 may be metered and is typically subject to transmission and distribution costs. In contrast, power consumed on the input side of AC-to-AC step-up transformer 760 may be considered behind-the-meter and is typically not subject to transmission and distribution costs. As such, one or more flexible datacenters 200 may be powered by three-phase solar-generated AC voltage 750. Specifically, in solar farm 700 applications, the three-phase behind-the-meter AC voltage used to power flexible datacenter 200 may be three-phase solar-generated AC voltage 750. As such, flexible datacenter 200 may reside behind-the-meter, avoid transmission and distribution costs, and may be dynamically powered when unutilized behind-the-meter power is available.

Unutilized behind-the-meter power availability may occur when there is excess local power generation. In high incident sunlight situations, solar farm 700 may generate more power than, for example, AC-to-AC step-up transformer 760 is rated for. In such situations, solar farm 700 may have to take steps to protect its equipment from damage, which may include taking one or more panels 710 offline or shunting their voltage to dummy loads or ground. Advantageously, one or more flexible datacenters 200 may be used to consume power on the input side of AC-to-AC step-up transformer 760, thereby allowing solar farm 700 to operate equipment within operating ranges while flexible datacenter 200 receives behind-the-meter power without transmission or distribution costs. The local station control system (not independently illustrated) of local station 775 may issue an operational directive to the one or more flexible datacenters 200 or to the remote master control system (420 of FIG. 4) to ramp-up to the desired power consumption level. When the operational directive requires the cooperative action of multiple flexible datacenters 200, the remote master control system (420 of FIG. 4) may determine how to

power each individual flexible datacenter 200 in accordance with the operational directive or provide an override to each flexible datacenter 200.

Another example of unutilized behind-the-meter power availability is when grid 790 cannot, for whatever reason, take the power being produced by solar farm 700. In such situations, solar farm 700 may have to take one or more panels 710 offline or shunt their voltage to dummy loads or ground. Advantageously, one or more flexible datacenters 200 may be used to consume power on the input side of AC-to-AC step-up transformer 760, thereby allowing solar farm 700 to either produce power to grid 790 at a lower level or shut down transformer 760 entirely while flexible datacenter 200 receives behind-the-meter power without transmission or distribution costs. The local station control system (not independently illustrated) of local station 775 or the grid operator (not independently illustrated) of grid 790 may issue an operational directive to the one or more flexible datacenters 200 or to the remote master control system (420 of FIG. 4) to ramp-up to the desired power consumption level. When the operational directive requires the cooperative action of multiple flexible datacenters 200, the remote master control system (420 of FIG. 4) may determine how to power each individual flexible datacenter 200 in accordance with the operational directive or provide an override to each flexible datacenter 200.

Another example of unutilized behind-the-meter power availability is when solar farm 700 is selling power to grid 790 at a negative price that is offset by a production tax credit. In certain circumstances, the value of the production tax credit may exceed the price solar farm 700 would have to pay to grid 790 to offload their generated power. Advantageously, one or more flexible datacenters 200 may be used to consume power behind-the-meter, thereby allowing solar farm 700 to produce and obtain the production tax credit, but sell less power to grid 790 at the negative price. The local station control system (not independently illustrated) of local station 775 may issue an operational directive to the one or more flexible datacenters 200 or to the remote master control system (420 of FIG. 4) to ramp-up to the desired power consumption level. When the operational directive requires the cooperative action of multiple flexible datacenter 200, the remote master control system (420 of FIG. 4) may determine how to power each individual flexible datacenter 200 in accordance with the operational directive or provide an override to each flexible datacenter 200.

Another example of unutilized behind-the-meter power availability is when solar farm 700 is selling power to grid 790 at a negative price because grid 790 is oversupplied or is instructed to stand down and stop producing altogether. The grid operator (not independently illustrated) may select certain power generation stations to go offline and stop producing power to grid 790. Advantageously, one or more flexible datacenters 200 may be used to consume power behind-the-meter, thereby allowing solar farm 700 to stop producing power to grid 790, but making productive use of the power generated behind-the-meter without transmission or distribution costs. The local station control system (not independently illustrated) of the local station 775 or the grid operator (not independently illustrated) of grid 790 may issue an operational directive to the one or more flexible datacenters 200 or to the remote master control system (420 of FIG. 4) to ramp-up to the desired power consumption level. When the operational directive requires the cooperative action of multiple flexible datacenters 200, the remote master control system (420 of FIG. 4) may determine how to power each individual flexible datacenter 200 in accordance with the operational directive or provide an override to each flexible datacenter 200.

dance with the operational directive or provide an override to each flexible datacenter 200.

Another example of unutilized behind-the-meter power availability is when solar farm 700 is producing power to grid 790 that is unstable, out of phase, or at the wrong frequency, or grid 790 is already unstable, out of phase, or at the wrong frequency for whatever reason. The grid operator (not independently illustrated) may select certain power generation stations to go offline and stop producing power to grid 790. Advantageously, one or more flexible datacenters 200 may be used to consume power behind-the-meter, thereby allowing solar farm 700 to stop producing power to grid 790, but make productive use of the power generated behind-the-meter without transmission or distribution costs. The local station control system (not independently illustrated) of local station 775 may issue an operational directive to the one or more flexible datacenters 200 or to the remote master control system (420 of FIG. 4) to ramp-up to the desired power consumption level. When the operational directive requires the cooperative action of multiple flexible datacenters 200, the remote master control system (420 of FIG. 4) may determine how to power each individual flexible datacenter 200 in accordance with the operational directive or provide an override to each flexible datacenter 200.

Further examples of unutilized behind-the-meter power availability is when solar farm 700 experiences intermittent cloud cover such that it is not economically feasible to power up certain components, such as, for example local station 775, but there may be sufficient behind-the-meter power availability to power one or more flexible datacenters 200. Similarly, unutilized behind-the-meter power availability may occur when solar farm 700 is starting up, or testing, one or more panels 710. Panels 710 are frequently offline for installation, maintenance, and service and must be tested prior to coming online as part of the array. One or more flexible datacenters 200 may be powered by one or more panels 710 that are offline from farm 700. The above-noted examples of when unutilized behind-the-meter power is available are merely exemplary and are not intended to limit the scope of what one of ordinary skill in the art would recognize as unutilized behind-the-meter power availability. Behind-the-meter power availability may occur anytime there is power available and accessible behind-the-meter that is not subject to transmission and distribution costs and there is an economic advantage to using it.

One of ordinary skill in the art will recognize that solar farm 700 and solar panel 710 may vary based on an application or design in accordance with one or more embodiments of the present invention.

FIG. 8 shows a flexible datacenter 200 powered by flare gas 800 in accordance with one or more embodiments of the present invention. Flare gas 800 is combustible gas produced as a product or by-product of petroleum refineries, chemical plants, natural gas processing plants, oil and gas drilling rigs, and oil and gas production facilities. Flare gas 800 is typically burned off through a flare stack (not shown) or vented into the air. In one or more embodiments of the present invention, flare gas 800 may be diverted 812 to a gas-powered generator that produces three-phase gas-generated AC voltage 822. This power may be considered behind-the-meter and is not subject to transmission and distribution costs. As such, one or more flexible datacenters 200 may be powered by three-phase gas-generated AC voltage. Specifically, the three-phase behind-the-meter AC voltage used to power flexible datacenter 200 may be three-phase gas-generated AC voltage 822. Accordingly,

flexible datacenter 200 may reside behind-the-meter, avoid transmission and distribution costs, and may be dynamically powered when unutilized behind-the-meter power is available.

FIG. 9 shows a method of dynamic power delivery to a flexible datacenter (200 of FIG. 2) using unutilized behind-the-meter power 900 in accordance with one or more embodiments of the present invention. In step 910, the datacenter control system (220 of FIG. 4), or the remote master control system (420 of FIG. 4), may monitor unutilized behind-the-meter power availability. In certain embodiments, monitoring may include receiving information or an operational directive from the local station control system (410 of FIG. 4) or the grid operator (440 of FIG. 4) corresponding to unutilized behind-the-meter power availability.

In step 920, the datacenter control system (220 of FIG. 4), or the remote master control system (420 of FIG. 4), may determine when a datacenter ramp-up condition is met. In certain embodiments, the datacenter ramp-up condition may be met when there is sufficient behind-the-meter power availability and there is no operational directive from the local station to go offline or reduce power. In step 930, the datacenter control system (220 of FIG. 4) may enable behind-the-meter power delivery to one or more computing systems (100 of FIG. 2). In step 940, once ramped-up, the datacenter control system (220 of FIG. 4) or the remote master control system (420 of FIG. 4) may direct one or more computing systems (100 of FIG. 2) to perform predetermined computational operations. In certain embodiments, the predetermined computational operations may include the execution of one or more hashing functions.

While operational, the datacenter control system (220 of FIG. 4), or the remote master control system (420 of FIG. 4), may receive an operational directive to modulate power consumption. In certain embodiments, the operational directive may be a directive to reduce power consumption. In such embodiments, the datacenter control system (220 of FIG. 4) or the remote master control system (420 of FIG. 4) may dynamically reduce power delivery to one or more computing systems (100 of FIG. 2) or dynamically reduce power consumption of one or more computing systems. In other embodiments, the operational directive may be a directive to provide a power factor correction factor. In such embodiments, the datacenter control system (220 of FIG. 4) or the remote master control system (420 of FIG. 4) may dynamically adjust power delivery to one or more computing systems (100 of FIG. 2) to achieve a desired power factor correction factor. In still other embodiments, the operational directive may be a directive to go offline or power down. In such embodiments, the datacenter control system (220 of FIG. 4) may disable power delivery to one or more computing systems (100 of FIG. 2).

The datacenter control system (220 of FIG. 4), or the remote master control system (420 of FIG. 4), may determine when a datacenter ramp-down condition is met. In certain embodiments, the datacenter ramp-down condition may be met if there is insufficient or anticipated to be insufficient behind-the-meter power availability or there is an operational directive from the local station to go offline or reduce power. The datacenter control system (220 of FIG. 4) may disable behind-the-meter power delivery to one or more computing systems (100 of FIG. 2). Once ramped-down, the datacenter control system (220 of FIG. 4) remains powered and in communication with the remote master

control system (420 of FIG. 4) so that it may dynamically power the flexible datacenter (200 of FIG. 2) when conditions change.

One of ordinary skill in the art will recognize that a datacenter control system (220 of FIG. 4) may dynamically modulate power delivery to one or more computing systems (100 of FIG. 2) of a flexible datacenter (200 of FIG. 2) based on unutilized behind-the-meter power availability or an operational directive. The flexible datacenter (200 of FIG. 2) may transition between a fully powered down state (while the datacenter control system remains powered), a fully powered up state, and various intermediate states in between. In addition, flexible datacenter (200 of FIG. 2) may have a blackout state, where all power consumption, including that of the datacenter control system (220 of FIG. 4) is halted. However, once the flexible datacenter (200 of FIG. 2) enters the blackout state, it will have to be manually rebooted to restore power to datacenter control system (220 of FIG. 4). Local station conditions or operational directives may cause flexible datacenter (200 of FIG. 2) to ramp-up, reduce power consumption, change power factor, or ramp-down.

Advantages of one or more embodiments of the present invention may include one or more of the following:

In one or more embodiments of the present invention, a method and system for dynamic power delivery to a flexible datacenter using unutilized energy sources provides a green solution to two prominent problems: the exponential increase in power required for growing blockchain operations and the unutilized and typically wasted energy generated from renewable energy sources.

In one or more embodiments of the present invention, a method and system for dynamic power delivery to a flexible datacenter using unutilized energy sources allows for the rapid deployment of mobile datacenters to local stations. The mobile datacenters may be deployed on site, near the source of power generation, and receive unutilized behind-the-meter power when it is available.

In one or more embodiments of the present invention, a method and system for dynamic power delivery to a flexible datacenter using unutilized energy sources allows for the power delivery to the datacenter to be modulated based on conditions or an operational directive received from the local station or the grid operator.

In one or more embodiments of the present invention, a method and system for dynamic power delivery to a flexible datacenter using unutilized energy sources may dynamically adjust power consumption by ramping-up, ramping-down, or adjusting the power consumption of one or more computing systems within the flexible datacenter.

In one or more embodiments of the present invention, a method and system for dynamic power delivery to a flexible datacenter using unutilized energy sources may be powered by unutilized behind-the-meter power that is free from transmission and distribution costs. As such, the flexible datacenter may perform computational operations, such as hashing function operations, with little to no energy cost.

In one or more embodiments of the present invention, a method and system for dynamic power delivery to a flexible datacenter using unutilized energy sources provides a number of benefits to the hosting local station. The local station may use the flexible datacenter to adjust a load, provide a power factor correction, to offload power, or operate in a manner that invokes a production tax credit.

While the present invention has been described with respect to the above-noted embodiments, those skilled in the art, having the benefit of this disclosure, will recognize that

US 11,016,456 B2

19

other embodiments may be devised that are within the scope of the invention as disclosed herein. Accordingly, the scope of the invention should be limited only by the appended claims.

What is claimed is:

1. A flexible datacenter comprising:
 - a behind-the-meter (BTM) power input system, wherein the BTM power input system is configured to receive power from a power generation station prior to the power undergoing step-up transformation for transmission to a power grid;
 - a power distribution system;
 - a plurality of computing systems; and
 - a datacenter control system configured to modulate power delivery to the plurality of computing systems based on an operational directive.
2. The flexible datacenter of claim 1, wherein the datacenter control system is configured to receive the operational directive from a remote master control system.
3. The flexible datacenter of claim 2, wherein the operational directive depends on BTM power availability at the flexible datacenter.
4. The flexible datacenter of claim 3, wherein the remote master control system is positioned remotely from the flexible datacenter.
5. The flexible datacenter of claim 3, wherein BTM power availability comprises:
 - one or more of excess local power generation at a local station level, local power generation subject to economic curtailment, local power generation subject to reliability, curtailment, local power generation subject to power factor correction, low local power generation, start up local power generation situations, transient local power generation situations, or testing local power generation situations where there is an economic advantage to using local BTM power generation to power the flexible datacenter.
6. The flexible datacenter of claim 3, wherein BTM power availability comprises:
 - excess local power generation at a local station level.
7. The flexible datacenter of claim 3, wherein BTM power availability comprises:
 - local power generation subject to economic curtailment.
8. The flexible datacenter of claim 3, wherein BTM power availability comprises:
 - local power generation subject to reliability curtailment.
9. The flexible datacenter of claim 3, wherein BTM power availability comprises:
 - local power generation c to power factor correction.
10. The flexible datacenter of claim 3, wherein BTM power availability comprises:
 - low local power generation.
11. The flexible datacenter of claim 3, wherein BTM power availability comprises:
 - start up local power generation situations.
12. The flexible datacenter of claim 3, wherein BTM power availability comprises:
 - transient local power generation situations.
13. The flexible datacenter of claim 3, wherein BTM power availability comprises:
 - testing local power generation situations where there is an economic advantage to using local BTM power generation to power the flexible datacenter.
14. The flexible datacenter of claim 1, wherein the datacenter control system is further configured to modulate power delivery to the plurality computing systems based on BTM power availability at the BTM power input system.

20

15. The flexible datacenter of claim 1, wherein the BTM power input system comprises:
 - an input configured to receive three-phase BTM AC voltage from the power generation station.
16. The flexible datacenter of claim 15, wherein the BTM power input system is configured to provide a single-phase AC voltage to the datacenter control system.
17. The flexible datacenter of claim 1, wherein the datacenter control system controllably enables or disables power delivery to respective computing systems of the plurality of computing systems.
18. The flexible datacenter of claim 1, wherein the datacenter control system is positioned remotely from the plurality of computing systems.
19. The flexible datacenter of claim 1, wherein the operational directive is a workload directive based on a current BTM power availability and a projected BTM power availability at the flexible datacenter.
20. The flexible datacenter of claim 1, wherein the datacenter control system is collocated with the plurality of computing systems.
21. The flexible datacenter of claim 1, wherein the datacenter control system is further configured to transfer a workload from the plurality of computing systems to a second plurality of computing systems, and wherein the second plurality of computing systems is positioned at a second flexible datacenter.
22. The flexible datacenter of claim 21, wherein the flexible datacenter is coupled to a first power generation station such that the plurality of computing systems operates using BTM power from the first power generation station; and
 - wherein the second flexible datacenter is coupled to a second power generation station such that the second plurality of computing systems operates using BTM power from the second power generation station.
23. The flexible datacenter of claim 22, wherein the datacenter control system is configured to transfer the workload from the plurality of computing systems to the second plurality of computing systems in response to detecting a change in BTM power availability at the plurality of computing systems.
24. The flexible datacenter of claim 23, wherein the datacenter control system is a remote master control system that is positioned remotely from the flexible datacenter and the second flexible datacenter.
25. A system comprising:
 - a first flexible datacenter that comprises:
 - a behind-the-meter (BTM) power input system, wherein the BTM power input system is configured to receive power from a power generation station prior to the power undergoing step-up transformation for transmission to a power grid;
 - a power distribution system; and
 - a plurality of computing systems configured to receive power from the BTM power input system via the power distribution system;
 - a second flexible datacenter; and
 - a routing control system configured to modulate power delivery to the first flexible datacenter and the second flexible datacenter.
26. The system of claim 25, wherein the routing control system is configured to transfer a workload from the plurality of computing systems to a second plurality of computing systems at the second flexible datacenter in response to detecting a change in BTM power availability at the first flexible datacenter.

US 11,016,456 B2

21

27. The system of claim 25, wherein the routing control system is positioned remotely from the first flexible datacenter and the second flexible datacenter, and wherein the routing control system is configured to modulate power delivery to the first flexible datacenter and the second flexible datacenter based on BTM power availability at the first flexible datacenter and the second flexible datacenter.

28. A method of dynamic power delivery to a flexible datacenter using behind-the-meter (BTM) power comprising:

monitoring, by a control system, BTM power availability at the flexible datacenter, wherein the flexible datacenter includes a BTM power input system configured to receive power from a power generation station prior to the power undergoing step-up transformation for transmission to a power grid;

determining that a ramp-up condition is met, wherein the ramp-up condition depends on BTM power availability at the flexible datacenter; and

22

based on determining that the ramp-up condition is met, providing an operational directive to a datacenter control system at the flexible datacenter, wherein the operational directive instructs the datacenter control system to enable BTM power delivery to one or more computing systems at the flexible datacenter.

29. The method of claim 28, further comprising: determining that a ramp-down condition is met, wherein the ramp-down condition depends on BTM power availability at the flexible datacenter; and

based on determining that the ramp-down condition is met, providing a second operational directive to the datacenter control system at the flexible datacenter, wherein the second operational directive instructs the datacenter control system to disable BTM power delivery to one or more computing systems at the flexible datacenter.

30. The method of claim 28, wherein the control system is positioned remotely from the flexible datacenter.

* * * * *

Exhibit 14

Redacted in its Entirety

Exhibit 15

Doc Code: TRACK1.REQ

Document Description: TrackOne Request

PTO/AIA/424 (04-14)

**CERTIFICATION AND REQUEST FOR PRIORITIZED EXAMINATION
UNDER 37 CFR 1.102(e) (Page 1 of 1)**

First Named Inventor:	McNamara	Nonprovisional Application Number (if known):	
Title of Invention:	Methods and Systems for Adjusting Power Consumption based on Fixed-Duration Power Option Agreement		

APPLICANT HEREBY CERTIFIES THE FOLLOWING AND REQUESTS PRIORITIZED EXAMINATION FOR THE ABOVE-IDENTIFIED APPLICATION.

- The processing fee set forth in 37 CFR 1.17(i)(1) and the prioritized examination fee set forth in 37 CFR 1.17(c) have been filed with the request. The publication fee requirement is met because that fee, set forth in 37 CFR 1.18(d), is currently \$0. The basic filing fee, search fee, and examination fee are filed with the request or have been already been paid. I understand that any required excess claims fees or application size fee must be paid for the application.
- I understand that the application may not contain, or be amended to contain, more than four independent claims, more than thirty total claims, or any multiple dependent claims, and that any request for an extension of time will cause an outstanding Track I request to be dismissed.
- The applicable box is checked below:

I. ☒ Original Application (Track One) - Prioritized Examination under § 1.102(e)(1)

- (a) The application is an original nonprovisional utility application filed under 35 U.S.C. 111(a). This certification and request is being filed with the utility application via EFS-Web.
---OR---
- (b) The application is an original nonprovisional plant application filed under 35 U.S.C. 111(a). This certification and request is being filed with the plant application in paper.
- An executed inventor's oath or declaration under 37 CFR 1.63 or 37 CFR 1.64 for each inventor, or the application data sheet meeting the conditions specified in 37 CFR 1.53(f)(3)(i) is filed with the application.

II. ☐ Request for Continued Examination - Prioritized Examination under § 1.102(e)(2)

- A request for continued examination has been filed with, or prior to, this form.
- If the application is a utility application, this certification and request is being filed via EFS-Web.
- The application is an original nonprovisional utility application filed under 35 U.S.C. 111(a), or is a national stage entry under 35 U.S.C. 371.
- This certification and request is being filed prior to the mailing of a first Office action responsive to the request for continued examination.
- No prior request for continued examination has been granted prioritized examination status under 37 CFR 1.102(e)(2).

Signature <u>/Alexander D. Georges/</u>	Date <u>December 4, 2019</u>
Name (Print/Typed) <u>Alexander D. Georges</u>	Practitioner Registration Number <u>70,534</u>
<p>Note: This form must be signed in accordance with 37 CFR 1.33. See 37 CFR 1.4(d) for signature requirements and certifications. Submit multiple forms if more than one signature is required.*</p>	
<p><input type="checkbox"/> *Total of _____ forms are submitted.</p>	

BB00000319

Privacy Act Statement

The **Privacy Act of 1974 (P.L. 93-579)** requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b)(2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether disclosure of these records is required by the Freedom of Information Act.
2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.
3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.
4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).
5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.
6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).
7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.
8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspection or an issued patent.
9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
16/702,931	12/04/2019	Michael T. McNamara	19-2071-US	2690
20306	7590	01/27/2020	EXAMINER	
MCDONNELL BOEHNEN HULBERT & BERGHOF LLP			EVERETT, CHRISTOPHER E	
300 S. WACKER DRIVE			ART UNIT	
32ND FLOOR			PAPER NUMBER	
CHICAGO, IL 60606			2116	
DATE MAILED: 01/27/2020				

Determination of Patent Term Adjustment under 35 U.S.C. 154 (b)
(Applications filed on or after May 29, 2000)

The Office has discontinued providing a Patent Term Adjustment (PTA) calculation with the Notice of Allowance.

Section 1(h)(2) of the AIA Technical Corrections Act amended 35 U.S.C. 154(b)(3)(B)(i) to eliminate the requirement that the Office provide a patent term adjustment determination with the notice of allowance. See Revisions to Patent Term Adjustment, 78 Fed. Reg. 19416, 19417 (Apr. 1, 2013). Therefore, the Office is no longer providing an initial patent term adjustment determination with the notice of allowance. The Office will continue to provide a patent term adjustment determination with the Issue Notification Letter that is mailed to applicant approximately three weeks prior to the issue date of the patent, and will include the patent term adjustment on the patent. Any request for reconsideration of the patent term adjustment determination (or reinstatement of patent term adjustment) should follow the process outlined in 37 CFR 1.705.

Any questions regarding the Patent Term Extension or Adjustment determination should be directed to the Office of Patent Legal Administration at (571)-272-7702. Questions relating to issue and publication fee payments should be directed to the Customer Service Center of the Office of Patent Publication at 1-(888)-786-0101 or (571)-272-4200.

OMB Clearance and PRA Burden Statement for PTOL-85 Part B

The Paperwork Reduction Act (PRA) of 1995 requires Federal agencies to obtain Office of Management and Budget approval before requesting most types of information from the public. When OMB approves an agency request to collect information from the public, OMB (i) provides a valid OMB Control Number and expiration date for the agency to display on the instrument that will be used to collect the information and (ii) requires the agency to inform the public about the OMB Control Number's legal significance in accordance with 5 CFR 1320.5(b).

The information collected by PTOL-85 Part B is required by 37 CFR 1.311. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.14. This collection is estimated to take 30 minutes to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, Virginia 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, Virginia 22313-1450. Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

Privacy Act Statement

The Privacy Act of 1974 (P.L. 93-579) requires that you be given certain information in connection with your submission of the attached form related to a patent application or patent. Accordingly, pursuant to the requirements of the Act, please be advised that: (1) the general authority for the collection of this information is 35 U.S.C. 2(b) (2); (2) furnishing of the information solicited is voluntary; and (3) the principal purpose for which the information is used by the U.S. Patent and Trademark Office is to process and/or examine your submission related to a patent application or patent. If you do not furnish the requested information, the U.S. Patent and Trademark Office may not be able to process and/or examine your submission, which may result in termination of proceedings or abandonment of the application or expiration of the patent.

The information provided by you in this form will be subject to the following routine uses:

1. The information on this form will be treated confidentially to the extent allowed under the Freedom of Information Act (5 U.S.C. 552) and the Privacy Act (5 U.S.C. 552a). Records from this system of records may be disclosed to the Department of Justice to determine whether disclosure of these records is required by the Freedom of Information Act.
2. A record from this system of records may be disclosed, as a routine use, in the course of presenting evidence to a court, magistrate, or administrative tribunal, including disclosures to opposing counsel in the course of settlement negotiations.
3. A record in this system of records may be disclosed, as a routine use, to a Member of Congress submitting a request involving an individual, to whom the record pertains, when the individual has requested assistance from the Member with respect to the subject matter of the record.
4. A record in this system of records may be disclosed, as a routine use, to a contractor of the Agency having need for the information in order to perform a contract. Recipients of information shall be required to comply with the requirements of the Privacy Act of 1974, as amended, pursuant to 5 U.S.C. 552a(m).
5. A record related to an International Application filed under the Patent Cooperation Treaty in this system of records may be disclosed, as a routine use, to the International Bureau of the World Intellectual Property Organization, pursuant to the Patent Cooperation Treaty.
6. A record in this system of records may be disclosed, as a routine use, to another federal agency for purposes of National Security review (35 U.S.C. 181) and for review pursuant to the Atomic Energy Act (42 U.S.C. 218(c)).
7. A record from this system of records may be disclosed, as a routine use, to the Administrator, General Services, or his/her designee, during an inspection of records conducted by GSA as part of that agency's responsibility to recommend improvements in records management practices and programs, under authority of 44 U.S.C. 2904 and 2906. Such disclosure shall be made in accordance with the GSA regulations governing inspection of records for this purpose, and any other relevant (i.e., GSA or Commerce) directive. Such disclosure shall not be used to make determinations about individuals.
8. A record from this system of records may be disclosed, as a routine use, to the public after either publication of the application pursuant to 35 U.S.C. 122(b) or issuance of a patent pursuant to 35 U.S.C. 151. Further, a record may be disclosed, subject to the limitations of 37 CFR 1.14, as a routine use, to the public if the record was filed in an application which became abandoned or in which the proceedings were terminated and which application is referenced by either a published application, an application open to public inspection or an issued patent.
9. A record from this system of records may be disclosed, as a routine use, to a Federal, State, or local law enforcement agency, if the USPTO becomes aware of a violation or potential violation of law or regulation.

Notice of Allowability	Application No. 16/702,931	Applicant(s) McNamara et al.	
	Examiner CHRISTOPHER E EVERETT	Art Unit 2116	AIA (FITF) Status Yes

– The MAILING DATE of this communication appears on the cover sheet with the correspondence address–

All claims being allowable, PROSECUTION ON THE MERITS IS (OR REMAINS) CLOSED in this application. If not included herewith (or previously mailed), a Notice of Allowance (PTOL-85) or other appropriate communication will be mailed in due course. **THIS NOTICE OF ALLOWABILITY IS NOT A GRANT OF PATENT RIGHTS.** This application is subject to withdrawal from issue at the initiative of the Office or upon petition by the applicant. See 37 CFR 1.313 and MPEP 1308.

1. ☒ This communication is responsive to 12/4/2019.
☐ A declaration(s)/affidavit(s) under **37 CFR 1.130(b)** was/were filed on ____.

2. ☐ An election was made by the applicant in response to a restriction requirement set forth during the interview on ____; the restriction requirement and election have been incorporated into this action.

3. ☒ The allowed claim(s) is/are 1-20. As a result of the allowed claim(s), you may be eligible to benefit from the **Patent Prosecution Highway** program at a participating intellectual property office for the corresponding application. For more information, please see http://www.uspto.gov/patents/init_events/pph/index.jsp or send an inquiry to PPHfeedback@uspto.gov.

4. ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

Certified copies:
a) ☐ All b) ☐ Some *c) ☐ None of the:
1. ☐ Certified copies of the priority documents have been received.
2. ☐ Certified copies of the priority documents have been received in Application No. _____.
3. ☐ Copies of the certified copies of the priority documents have been received in this national stage application from the International Bureau (PCT Rule 17.2(a)).

* Certified copies not received: _____.

Applicant has **THREE MONTHS FROM THE "MAILING DATE"** of this communication to file a reply complying with the requirements noted below. Failure to timely comply will result in **ABANDONMENT** of this application.
THIS THREE-MONTH PERIOD IS NOT EXTENDABLE.

5. ☐ **CORRECTED DRAWINGS** (as "replacement sheets") must be submitted.
☐ including changes required by the attached Examiner's Amendment / Comment or in the Office action of Paper No./Mail Date _____.

Identifying indicia such as the application number (see 37 CFR 1.84(c)) should be written on the drawings in the front (not the back) of each sheet. Replacement sheet(s) should be labeled as such in the header according to 37 CFR 1.121(d).

6. ☐ **DEPOSIT OF and/or INFORMATION** about the deposit of **BIOLOGICAL MATERIAL** must be submitted. Note the attached Examiner's comment regarding **REQUIREMENT FOR THE DEPOSIT OF BIOLOGICAL MATERIAL**.

Attachment(s)

1. <input checked="" type="checkbox"/> Notice of References Cited (PTO-892) 2. <input checked="" type="checkbox"/> Information Disclosure Statements (PTO/SB/08), Paper No./Mail Date <u>12/17/2019</u> . 3. <input type="checkbox"/> Examiner's Comment Regarding Requirement for Deposit of Biological Material _____. 4. <input type="checkbox"/> Interview Summary (PTO-413), Paper No./Mail Date _____.	5. <input type="checkbox"/> Examiner's Amendment/Comment 6. <input checked="" type="checkbox"/> Examiner's Statement of Reasons for Allowance 7. <input type="checkbox"/> Other _____.
--	--

/Christopher E. Everett/
Primary Examiner, Art Unit 2116

Application/Control Number: 16/702,931
Art Unit: 2116

Page 2

DETAILED ACTION

1. The Office Action is responsive to the communication filed on 12/4/2019.
2. Claims 1-20 are pending.

Notice of Pre-AIA or AIA Status

3. The present application, filed on or after March 16, 2013, is being examined under the first inventor to file provisions of the AIA.

Allowable Subject Matter

4. Claims 1-20 are allowed.
5. The following is an examiner's statement of reasons for allowance:

Claims 1-16

Regarding claim 1, the prior art as described in the prosecution history describes:

A system comprising:

- a set of computing systems, wherein the set of computing systems is configured to perform computational operations using power from a power grid;
- a control system configured to:
 - monitor a set of conditions;

Application/Control Number: 16/702,931
Art Unit: 2116

Page 3

However, regarding claim 1, the prior art as described in the prosecution history does not describe:

- receive power option data based, at least in part, on a power option agreement, wherein the power option data specify:
 - (i) a set of minimum power thresholds, and
 - (ii) a set of time intervals,
 - wherein each minimum power threshold in the set of minimum power thresholds is associated with a time interval in the set of time intervals;
- responsive to receiving the power option data, determine a performance strategy for the set of computing systems based on a combination of at least a portion of the power option data and at least one condition in the set of conditions,
 - wherein the performance strategy comprises a power consumption target for the set of computing systems for each time interval in the set of time intervals,
 - wherein each power consumption target is equal to or greater than the minimum power threshold associated with each time interval;and
- provide instructions to the set of computing systems to perform one or more computational operations based on the performance strategy.

Application/Control Number: 16/702,931
Art Unit: 2116

Page 4

Dependent claims 2-16 depend from independent claim 1 and are allowable for the same reasons as described above.

Claims 17-19

Independent claim 17 is substantially similar to independent claim 1 and is allowable for the same reasons as outlined above with respect to claim 1. Dependent claims 18-19 depend from independent claim 17 and are allowable for the same reasons as described above.

Claim 20

Independent claim 20 is substantially similar to independent claim 1 and is allowable for the same reasons as outlined above with respect to claim 1.

6. Any comments considered necessary by applicant must be submitted no later than the payment of the issue fee and, to avoid processing delays, should preferably accompany the issue fee. Such submissions should be clearly labeled “Comments on Statement of Reasons for Allowance.”

Application/Control Number: 16/702,931
Art Unit: 2116

Page 5

Conclusion

Any inquiry concerning this communication or earlier communications from the examiner should be directed to CHRISTOPHER E EVERETT whose telephone number is (571)272-2851. The examiner can normally be reached on Monday-Friday 8:00 am to 5:00 pm (Eastern).

Examiner interviews are available via telephone, in-person, and video conferencing using a USPTO supplied web-based collaboration tool. To schedule an interview, applicant is encouraged to use the USPTO Automated Interview Request (AIR) at <http://www.uspto.gov/interviewpractice>.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Kenneth Lo can be reached on 571-272-9774. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

BB00000561

Application/Control Number: 16/702,931
Art Unit: 2116

Page 6

/Christopher E. Everett/
Primary Examiner, Art Unit 2116

Notice of References Cited	Application/Control No. 16/702,931	Applicant(s)/Patent Under Reexamination McNamara et al.	
	Examiner CHRISTOPHER E EVERETT	Art Unit 2116	Page 1 of 1

U.S. PATENT DOCUMENTS

*		Document Number Country Code-Number-Kind Code	Date MM-YYYY	Name	CPC Classification	US Classification
*	A	US-20160187906-A1	06-2016	Bodas; Devadatta V.	G05B15/02	700/287
*	B	US-20090070611-A1	03-2009	Bower, III; Fred A.	G06F1/3203	713/322
*	C	US-20090089595-A1	04-2009	Brey; Thomas M.	G06F1/3203	713/300
*	D	US-20130086404-A1	04-2013	Sankar; Sriram	G06F1/305	713/324
*	E	US-20050203761-A1	09-2005	Barr, Andrew Harvey	G06F1/26	713/320
*	F	US-8700929-B1	04-2014	Weber; Wolf-Dietrich	G06F30/13	713/310
	G					
	H					
	I					
	J					
	K					
	L					
	M					


FOREIGN PATENT DOCUMENTS

*		Document Number Country Code-Number-Kind Code	Date MM-YYYY	Country	Name	CPC Classification
	N					
	O					
	P					
	Q					
	R					
	S					
	T					

NON-PATENT DOCUMENTS

*		Include as applicable: Author, Title Date, Publisher, Edition or Volume, Pertinent Pages)
	U	
	V	
	W	
	X	


*A copy of this reference is not being furnished with this Office action. (See MPEP § 707.05(a).)
Dates in MM-YYYY format are publication dates. Classifications may be US or foreign.

Issue Classification 	Application/Control No. 16/702,931	Applicant(s)/Patent Under Reexamination McNamara et al.
	Examiner CHRISTOPHER E EVERETT	Art Unit 2116

CPC							Type		Version
Symbol									
H02J	/	3	/	14			F		2013-01-01
H02J	/	3	/	008			I		2013-01-01
G06F	/	1	/	3203			I		2013-01-01

CPC Combination Sets										Type		Set	Ranking	Version
Symbol														
	/		/											

NONE		Total Claims Allowed:	
(Assistant Examiner)	(Date)	20	
/Christopher E. Everett/ Primary Examiner, Art Unit 2116	24 January 2020	O.G. Print Claim(s)	O.G. Print Figure
(Primary Examiner)	(Date)	1	12

Issue Classification 	Application/Control No. 16/702,931	Applicant(s)/Patent Under Reexamination McNamara et al.
	Examiner CHRISTOPHER E EVERETT	Art Unit 2116


INTERNATIONAL CLASSIFICATION			
CLAIMED			
H02J3/14	/	3	/ 14

NON-CLAIMED			
/		/	

US ORIGINAL CLASSIFICATION	
CLASS	SUBCLASS


CROSS REFERENCES(S)					
CLASS	SUBCLASS (ONE SUBCLASS PER BLOCK)				

NONE		Total Claims Allowed:	
(Assistant Examiner)	(Date)	20	
/Christopher E. Everett/ Primary Examiner, Art Unit 2116	24 January 2020	O.G. Print Claim(s)	O.G. Print Figure
(Primary Examiner)	(Date)	1	12

Issue Classification 	Application/Control No. 16/702,931	Applicant(s)/Patent Under Reexamination McNamara et al.
	Examiner CHRISTOPHER E EVERETT	Art Unit 2116

<input checked="" type="checkbox"/> Claims renumbered in the same order as presented by applicant <input type="checkbox"/> CPA <input type="checkbox"/> T.D. <input type="checkbox"/> R.1.47															
CLAIMS															
Final	Original	Final	Original	Final	Original	Final	Original	Final	Original	Final	Original	Final	Original	Final	Original

NONE (Assistant Examiner) _____ (Date) _____		Total Claims Allowed: 20	
/Christopher E. Everett/ Primary Examiner, Art Unit 2116 (Primary Examiner) _____ (Date) _____		24 January 2020 (Date)	O.G. Print Claim(s) 1 O.G. Print Figure 12

Search Notes 	Application/Control No. 16/702,931	Applicant(s)/Patent Under Reexamination McNamara et al.
	Examiner CHRISTOPHER E EVERETT	Art Unit 2116

CPC - Searched*		
Symbol	Date	Examiner
(H02J3/14 or G06F1/3203 or H02J3/008)	01/24/2020	CEE


CPC Combination Sets - Searched*		
Symbol	Date	Examiner

US Classification - Searched*			
Class	Subclass	Date	Examiner

* See search history printout included with this form or the SEARCH NOTES box below to determine the scope of the search.

Search Notes		
Search Notes	Date	Examiner
Search of EAST using inventor, applicant, and text searches	01/24/2020	CEE
Search of (H02J3/14 or G06F1/3203 or H02J3/008) using text searches	01/24/2020	CEE
Search of 700/295 using text searches	01/24/2020	CEE
Search of Google patent using text searches	01/24/2020	CEE
Search of Google scholar using text searches	01/24/2020	CEE
Search of IP.com using text searches	01/22/2020	CEE
Search of STIC NPL database using text searches	01/22/2020	CEE
Search of DAV using inventor searches	01/24/2020	CEE

/Christopher E. Everett/ Primary Examiner, Art Unit 2116	
---	--

<i>Search Notes</i> 	Application/Control No. 16/702,931	Applicant(s)/Patent Under Reexamination McNamara et al.
	Examiner CHRISTOPHER E EVERETT	Art Unit 2116

Interference Search			
US Class/CPC Symbol	US Subclass/CPC Group	Date	Examiner
	Search of USPAT and USPGPUB using text searches	01/24/2020	CEE

/Christopher E. Everett/ Primary Examiner, Art Unit 2116	
---	--

AO 120 (Rev. 08/10)

TO: Mail Stop 8 Director of the U.S. Patent and Trademark Office P.O. Box 1450 Alexandria, VA 22313-1450	REPORT ON THE FILING OR DETERMINATION OF AN ACTION REGARDING A PATENT OR TRADEMARK
---	---

In Compliance with 35 U.S.C. § 290 and/or 15 U.S.C. § 1116 you are hereby advised that a court action has been filed in the U.S. District Court Western District of Texas, Waco Division on the following

☐ Trademarks or ☒ Patents. (☐ the patent action involves 35 U.S.C. § 292.);

DOCKET NO.	DATE FILED 8/14/2020	U.S. DISTRICT COURT Western District of Texas, Waco Division
PLAINTIFF Lancium LLC		DEFENDANT Layer1 Technologies, Inc.
PATENT OR TRADEMARK NO.	DATE OF PATENT OR TRADEMARK	HOLDER OF PATENT OR TRADEMARK
1 10,608,433	3/31/2020	Lancium LLC
2		
3		
4		
5		

In the above—entitled case, the following patent(s)/ trademark(s) have been included:

DATE INCLUDED	INCLUDED BY <input type="checkbox"/> Amendment <input type="checkbox"/> Answer <input type="checkbox"/> Cross Bill <input type="checkbox"/> Other Pleading	
PATENT OR TRADEMARK NO.	DATE OF PATENT OR TRADEMARK	HOLDER OF PATENT OR TRADEMARK
1		
2		
3		
4		
5		

In the above—entitled case, the following decision has been rendered or judgement issued:

DECISION/JUDGEMENT		
CLERK	(BY) DEPUTY CLERK	DATE

Copy 1—Upon initiation of action, mail this copy to Director Copy 3—Upon termination of action, mail this copy to Director
 Copy 2—Upon filing document adding patent(s), mail this copy to Director Copy 4—Case file copy

BB00000667

Exhibit 16

Redacted in its Entirety

Exhibit 17

US010608433B1

(12) **United States Patent**
McNamara et al.

(10) **Patent No.:** **US 10,608,433 B1**
(45) **Date of Patent:** **Mar. 31, 2020**

(54) **METHODS AND SYSTEMS FOR ADJUSTING POWER CONSUMPTION BASED ON A FIXED-DURATION POWER OPTION AGREEMENT**

7,143,300 B2 11/2006 Potter et al.
7,647,516 B2 1/2010 Ranganathan et al.
(Continued)

FOREIGN PATENT DOCUMENTS

(71) Applicant: **Lancium LLC**, Houston, TX (US)

CN 103163904 A 6/2013
KR 20090012523 A 2/2009
WO 2015199629 A1 12/2015

(72) Inventors: **Michael T. McNamara**, Newport Beach, CA (US); **Raymond E. Cline, Jr.**, Houston, TX (US)

OTHER PUBLICATIONS

(73) Assignee: **Lancium LLC**, Houston, TX (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

Bird et al., "Wind and Solar Energy Curtailment: Experience and Practices in the United States," National Renewable Energy Lab (NREL), Technical Report NREL/TP-6A20-60983, Mar. 2014, 58 pages.

(Continued)

(21) Appl. No.: **16/702,931**

(22) Filed: **Dec. 4, 2019**

Primary Examiner — Christopher E. Everett

(74) *Attorney, Agent, or Firm* — McDonnell Boehnen Hulbert & Berghoff LLP

Related U.S. Application Data

(60) Provisional application No. 62/927,119, filed on Oct. 28, 2019.

(51) **Int. Cl.**
H02J 3/14 (2006.01)
H02J 3/00 (2006.01)
G06F 1/3203 (2019.01)

(52) **U.S. Cl.**
CPC **H02J 3/14** (2013.01); **G06F 1/3203** (2013.01); **H02J 3/008** (2013.01)

(58) **Field of Classification Search**
CPC H02J 3/14; H02J 3/008; G06F 1/3203
See application file for complete search history.

(57) **ABSTRACT**


Examples relate to adjusting load power consumption based on a power option agreement. A computing system may receive power option data that is based on a power option agreement and specify minimum power thresholds associated with time intervals. The computing system may determine a performance strategy for a load (e.g., set of computing systems) based on a combination of the power option data and one or more monitored conditions. The performance strategy may specify a power consumption target for the load for each time interval such that each power consumption target is equal to or greater than the minimum power threshold associated with each time interval. The computing system may provide instructions the set of computing systems to perform one or more computational operations based on the performance strategy.

(56) **References Cited**

U.S. PATENT DOCUMENTS

6,288,456 B1 9/2001 Cratty
6,633,823 B2 10/2003 Bartone et al.

20 Claims, 16 Drawing Sheets



US 10,608,433 B1

Page 2

(56)

References Cited

U.S. PATENT DOCUMENTS

7,702,931	B2	4/2010	Goodrum et al.	
7,779,276	B2	8/2010	Bolan et al.	
7,861,102	B1	12/2010	Ranganathan et al.	
7,921,315	B2	4/2011	Langgood et al.	
7,970,561	B2	6/2011	Pfeiffer	
8,001,403	B2	8/2011	Hamilton et al.	
8,006,108	B2	8/2011	Brey et al.	
8,214,843	B2	7/2012	Boss et al.	
8,374,928	B2	2/2013	Gopisetty et al.	
8,447,993	B2	5/2013	Greene et al.	
8,571,820	B2	10/2013	Pfeiffer	
8,627,123	B2	1/2014	Jain et al.	
8,700,929	B1*	4/2014	Weber	G06F 30/13 713/310
8,789,061	B2	7/2014	Pavel et al.	
8,799,690	B2	8/2014	Dawson et al.	
9,003,211	B2	4/2015	Pfeiffer	
9,003,216	B2	4/2015	Sankar et al.	
9,026,814	B2	5/2015	Aasheim et al.	
9,207,993	B2	12/2015	Jain	
9,218,035	B2	12/2015	Li et al.	
9,552,234	B2	1/2017	Boldyrev et al.	
10,367,353	B1	7/2019	McNamara et al.	
10,367,535	B2	7/2019	Corse et al.	
10,444,818	B1	10/2019	McNamara et al.	
10,452,127	B1	10/2019	McNamara et al.	
10,497,072	B2	12/2019	Hooshmand et al.	
2002/0072868	A1	6/2002	Bartone et al.	
2003/0074464	A1	4/2003	Bohrer et al.	
2005/0203761	A1*	9/2005	Barr	G06F 1/26 713/320
2006/0161765	A1	7/2006	Cromer et al.	
2008/0030078	A1	2/2008	Whitted et al.	
2008/0094797	A1	4/2008	Coglitore et al.	
2009/0055665	A1	2/2009	Maglione et al.	
2009/0070611	A1*	3/2009	Bower, III	G06F 1/3203 713/322
2009/0089595	A1*	4/2009	Brey	G06F 1/3203 713/300
2010/0211810	A1	8/2010	Zacho	
2010/0328849	A1	12/2010	Ewing et al.	
2011/0072289	A1	3/2011	Kato	
2012/0000121	A1	1/2012	Swann	
2012/0072745	A1	3/2012	Ahluwalia et al.	
2012/0300524	A1	11/2012	Fornage et al.	
2012/0324259	A1	12/2012	Aasheim et al.	
2013/0006401	A1	1/2013	Shan	
2013/0063991	A1	3/2013	Xiao et al.	
2013/0086404	A1*	4/2013	Sankar	G06F 1/305 713/324
2013/0187464	A1	7/2013	Smith et al.	
2013/0227139	A1	8/2013	Suffling	
2013/0306276	A1	11/2013	Duchesneau	
2014/0137468	A1	5/2014	Ching	
2014/0379156	A1	12/2014	Kamel et al.	
2015/0121113	A1	4/2015	Ramamurthy et al.	
2015/0155712	A1	6/2015	Mondal	
2015/0229227	A1	8/2015	Aeloiza et al.	
2015/0277410	A1	10/2015	Gupta et al.	
2015/0278968	A1	10/2015	Steven et al.	
2016/0170469	A1	6/2016	Schgal et al.	
2016/0172900	A1	6/2016	Welch, Jr.	
2016/0187906	A1*	6/2016	Bodas	G05B 15/02 700/287

2016/0198656	A1	7/2016	McNamara et al.
2016/0212954	A1	7/2016	Argento
2016/0324077	A1	11/2016	Frantzen et al.
2017/0023969	A1	1/2017	Shows et al.
2017/0104336	A1	4/2017	Elbsat et al.
2017/0261949	A1	9/2017	Hoffmann et al.
2018/0144414	A1	5/2018	Lee et al.
2018/0202825	A1	7/2018	You et al.
2018/0240112	A1	8/2018	Castinado et al.
2018/0366978	A1	12/2018	Matan et al.
2018/0367320	A1	12/2018	Montalvo
2019/0052094	A1	2/2019	Pmsvsv et al.
2019/0168630	A1	6/2019	Mrlik et al.
2019/0258307	A1	8/2019	Shaikh et al.
2019/0280521	A1	9/2019	Lundstrom et al.
2019/0318327	A1	10/2019	Sowell et al.
2019/0324820	A1	10/2019	Krishnan et al.

OTHER PUBLICATIONS

EPEX Spot, "How They Occur, What They Mean," https://www.epexspot.com/en/company-info/basics_of_the_power_market/negative_prices, 2018, 2 pages.

Final Office Action dated Oct. 1, 2019 for U.S. Appl. No. 16/175,246, filed Oct. 30, 2018, 18 pages.

Ghamkhari et al., "Optimal Integration of Renewable Energy Resources in Data Centers with Behind-the-Meter Renewable Generator," Department of Electrical and Computer Engineering Texas Tech University, 2012, pp. 3340-3444.

International Search Report and Written Opinion of PCT Application No. PCT/US2018/017955, dated Apr. 30, 2018, 22 pages.

International Search Report and Written Opinion of PCT Application No. PCT/US2018/017950, dated May 31, 2018, 15 pages.

Non-Final Office Action dated Dec. 5, 2019 for U.S. Appl. No. 16/529,360, filed Aug. 1, 2019, 72 pages.

Non-Final Office Action dated Dec. 10, 2019 for U.S. Appl. No. 16/596,190, filed Oct. 8, 2019, 72 pages.

Non-Final Office Action dated Nov. 14, 2019 for U.S. Appl. No. 16/132,098, filed Sep. 14, 2018, 25 pages.

Non-Final Office Action dated Nov. 21, 2019 for U.S. Appl. No. 16/529,402, filed Aug. 1, 2019, 57 pages.

Non-Final Office Action dated Dec. 11, 2019 on for U.S. Appl. No. 16/132,062, filed Sep. 14, 2018, 17 pages.

Non-Final Office Action dated Dec. 10, 2019 for U.S. Appl. No. 16/528,348, filed Oct. 8, 2019, 33 pages.

Notice of Allowance dated Apr. 2, 2019, for U.S. Appl. No. 16/175,335, filed Oct. 30, 2018, 12 pages.

Notice of Allowance dated Aug. 15, 2019, for U.S. Appl. No. 16/175,146, filed Oct. 30, 2018, 17 pages.

Notice of Allowance dated Jul. 29, 2019, for U.S. Appl. No. 16/245,532, filed Jan. 11, 2019, 13 pages.

Rahimi, Farrokh, "Using a Transactive Energy Framework," IEEE Electrification Magazine, Dec. 2016, pp. 23-29.

Soluna, "Powering the Block Chain," Aug. 2018, version 1.1, 29 pages.

Wilson, Joseph Nathanael, "A Utility-Scale Deployment Project of Behind-the-Meter Energy Storage for Use in Ancillary Services, Energy Resiliency, Grid Infrastructure Investment Deferral, and Demand-Response Integration," Portland State University, 2016, 154 pages.


* cited by examiner

U.S. Patent

Mar. 31, 2020

Sheet 1 of 16

US 10,608,433 B1



PRIOR ART
FIGURE 1

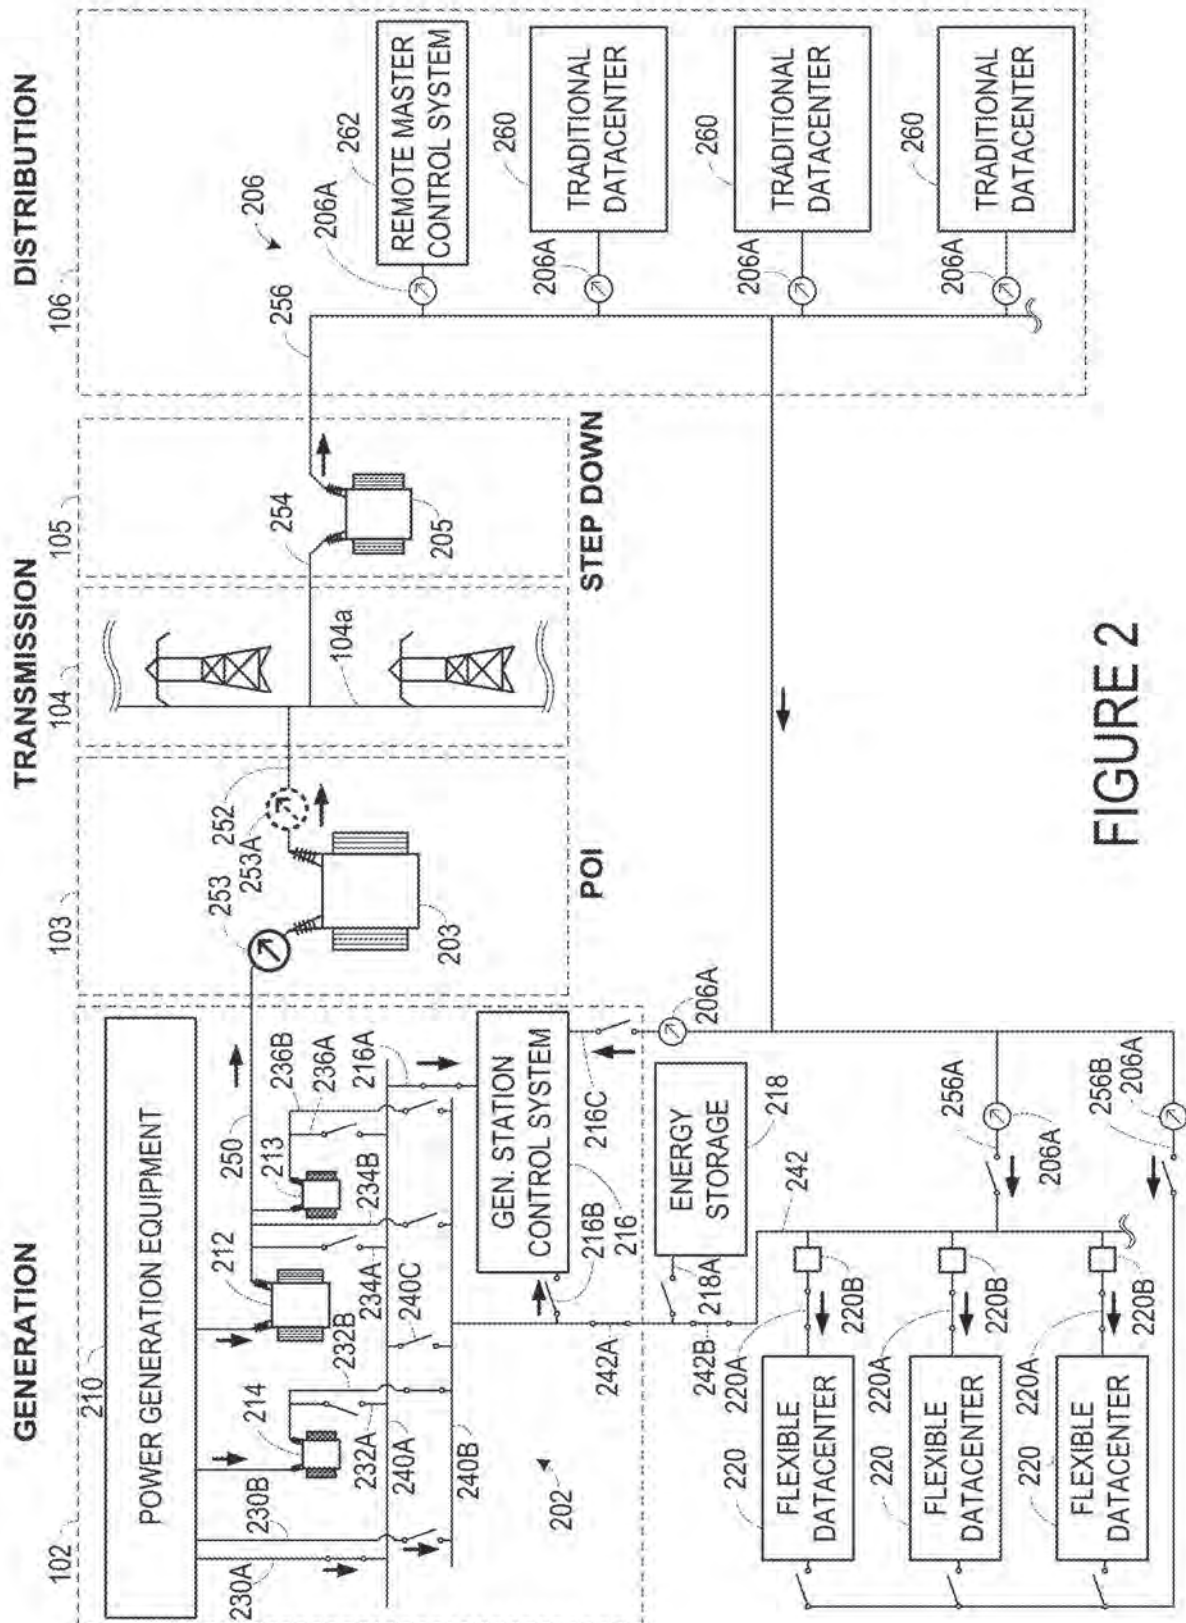


FIGURE 2

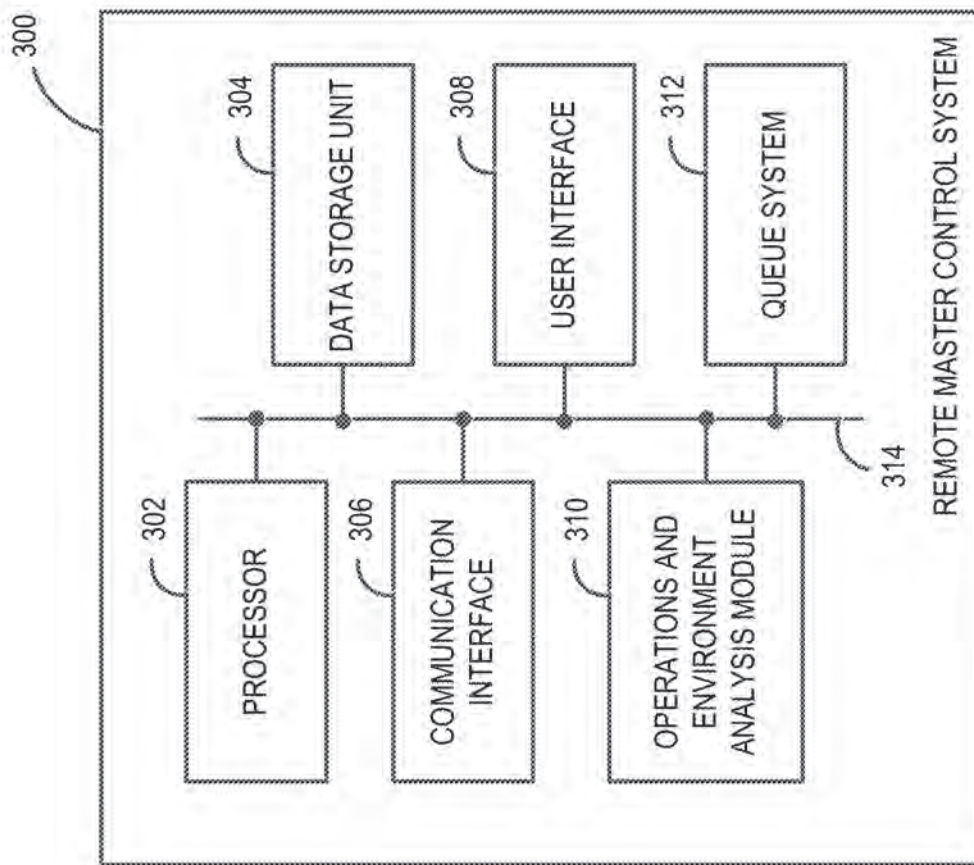


FIGURE 3

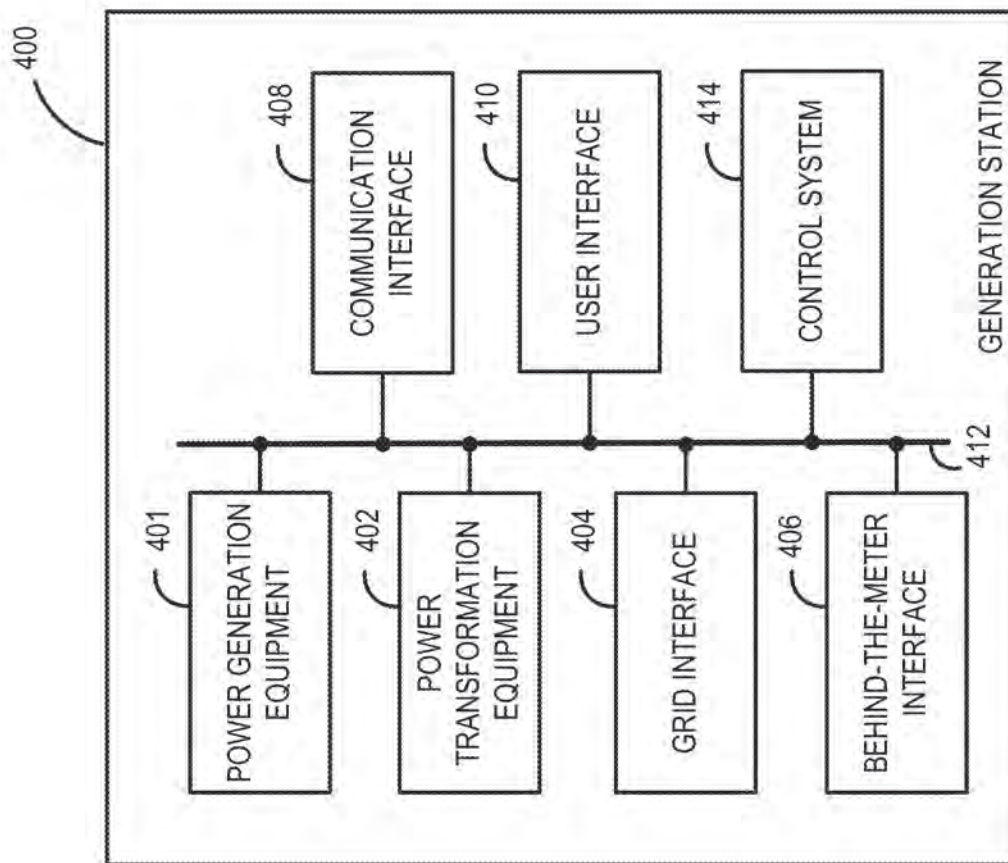


FIGURE 4





FIGURE 5




U.S. Patent

Mar. 31, 2020

Sheet 7 of 16

US 10,608,433 B1







FIGURE 7






FIGURE 9

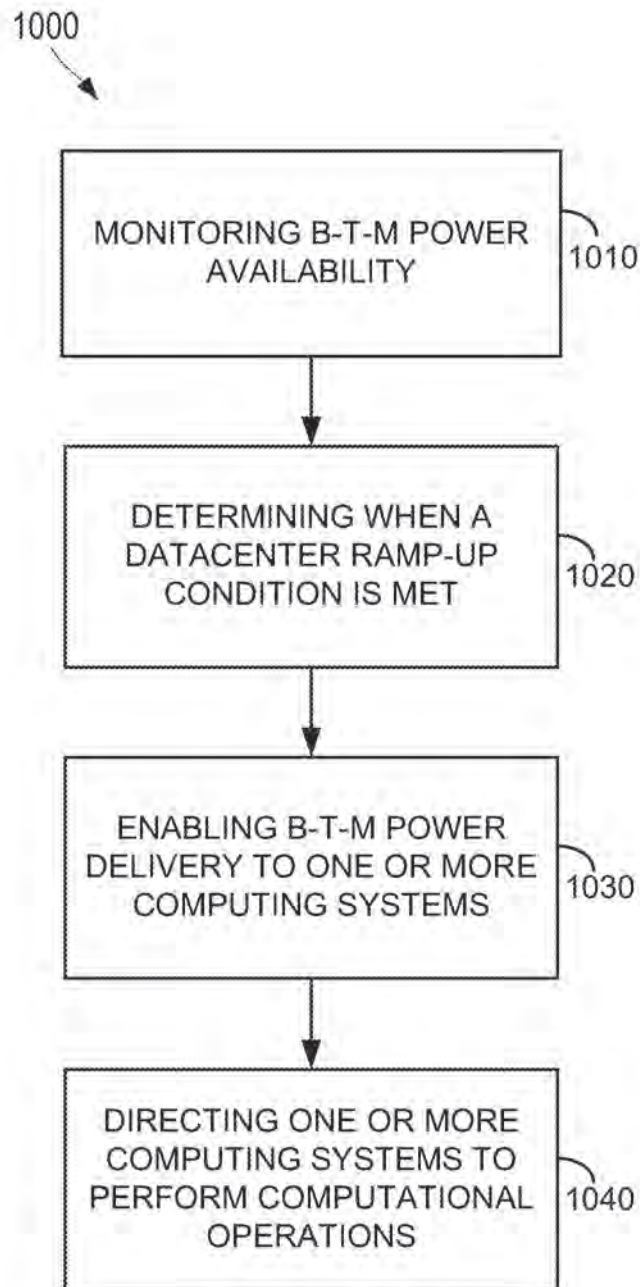


FIGURE 10A

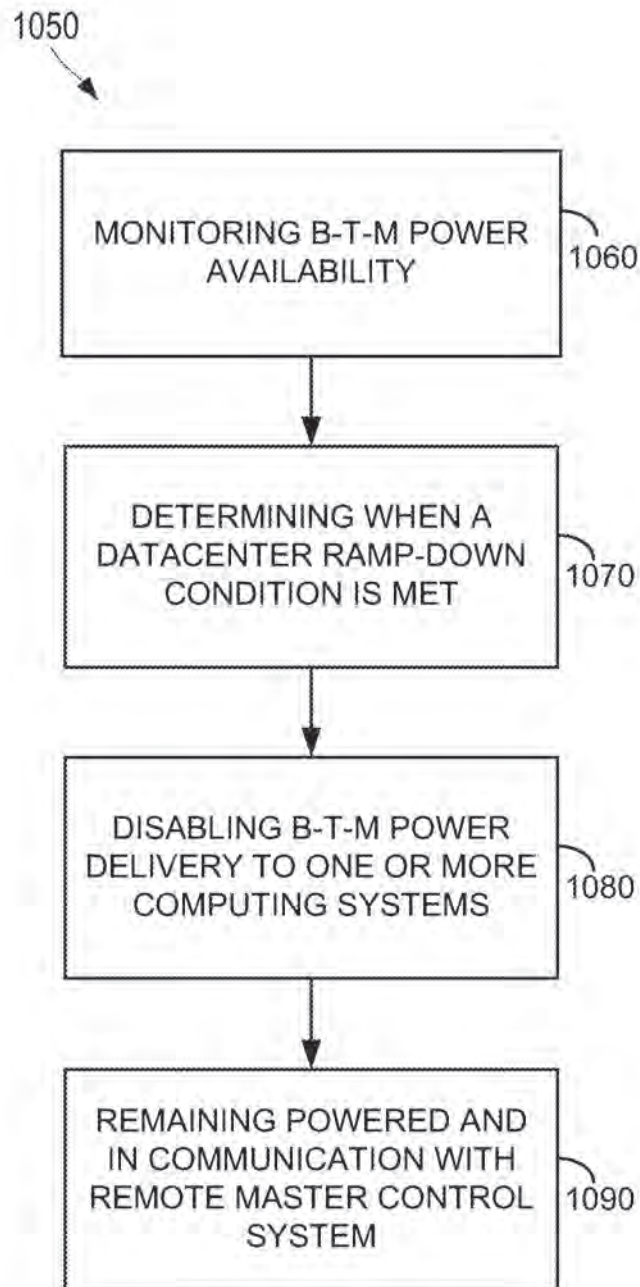


FIGURE 10B




FIGURE 11




FIGURE 12

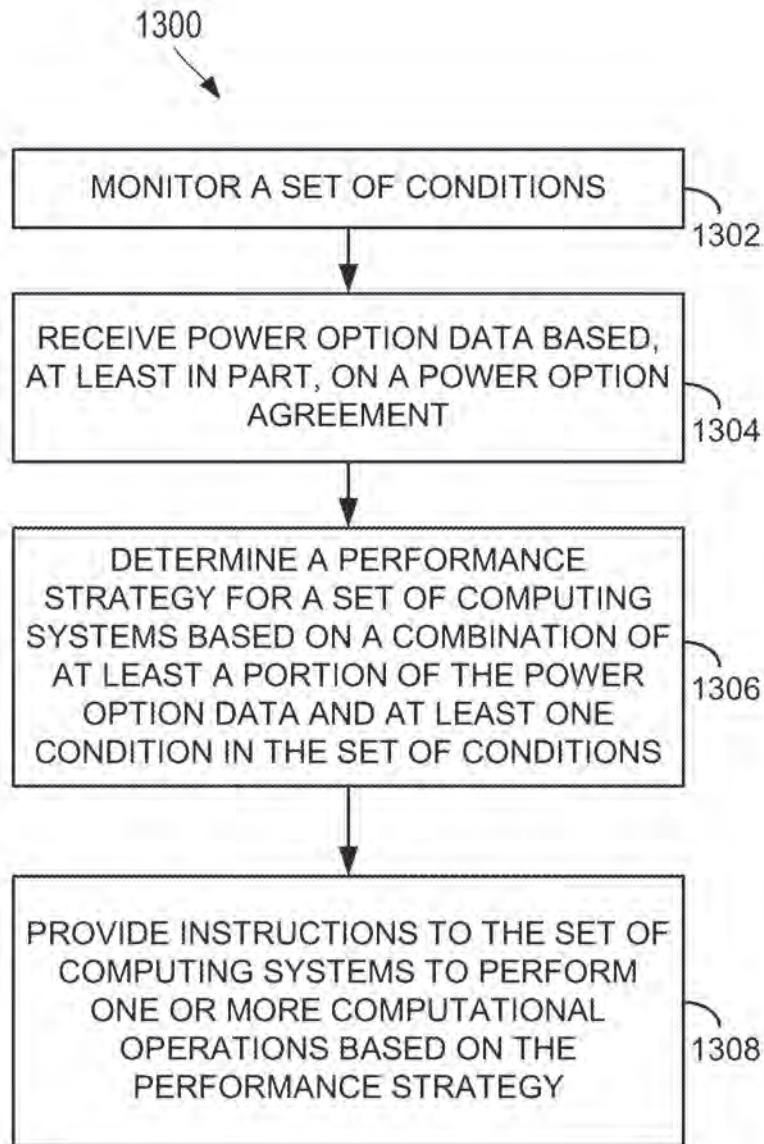


FIGURE 13

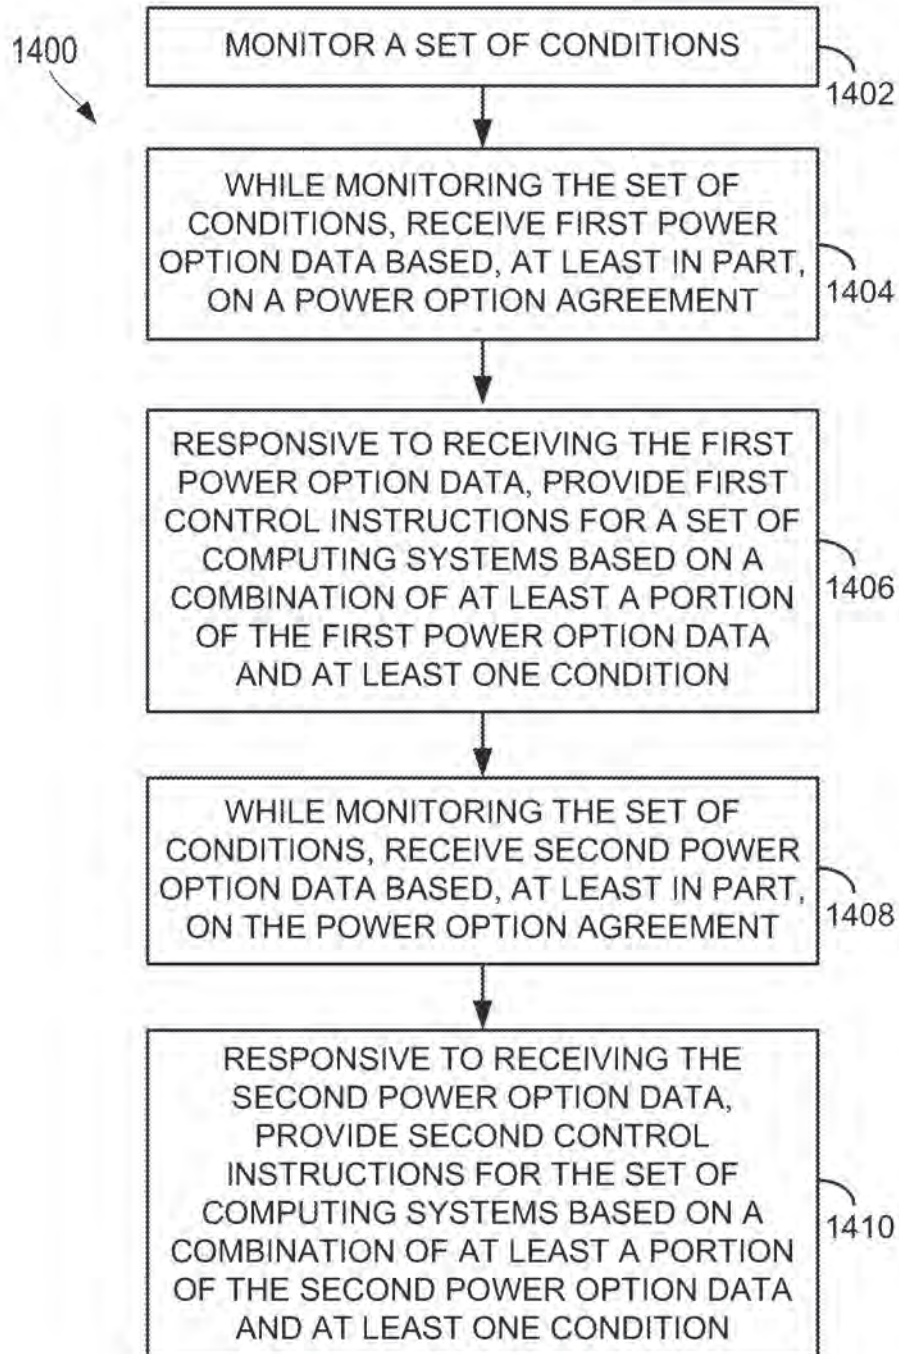


FIGURE 14

US 10,608,433 B1

1

METHODS AND SYSTEMS FOR ADJUSTING POWER CONSUMPTION BASED ON A FIXED-DURATION POWER OPTION AGREEMENT

CROSS-REFERENCE TO RELATED APPLICATIONS

The present application claims priority to U.S. Provisional Patent Application No. 62/927,119, filed Oct. 28, 2019, the entire contents of which are herein incorporated by reference.

FIELD

This specification relates to power consumption adjustments when using grid power and/or intermittent behind-the-meter power.

BACKGROUND

"Electrical grid" or "grid," as used herein, refers to a Wide Area Synchronous Grid (also known as an Interconnection), and is a regional scale or greater electric power grid that operates at a synchronized frequency and is electrically tied together during normal system conditions. An electrical grid delivers electricity from generation stations to consumers. An electrical grid includes: (i) generation stations that produce electrical power at large scales for delivery through the grid, (ii) high voltage transmission lines that carry that power from the generation stations to demand centers, and (iii) distribution networks carry that power to individual customers.

FIG. 1 illustrates a typical electrical grid, such as a North American Interconnection or the synchronous grid of Continental Europe (formerly known as the UCTE grid). The electrical grid of FIG. 1 can be described with respect to the various segments that make up the grid.

A generation segment **102** includes one or more generation stations that produce utility-scale electricity (typically >50 MW), such as a nuclear plant **102a**, a coal plant **102b**, a wind power station (i.e., wind farm) **102c**, and/or a photovoltaic power station (i.e., a solar farm) **102d**. Generation stations are differentiated from building-mounted and other decentralized or local wind or solar power applications because they supply power at the utility level and scale (>50 MW), rather than to a local user or users. The primary purpose of generation stations is to produce power for distribution through the grid, and in exchange for payment for the supplied electricity. Each of the generation stations **102a-d** includes power generation equipment **102e-h**, respectively, typically capable of supply utility-scale power (>50 MW). For example, the power generation equipment **102g** at wind power station **102c** includes wind turbines, and the power generation equipment **102h** at photovoltaic power station **102d** includes photovoltaic panels.

Each of the generation stations **102a-d** may further include station electrical equipment **102i-1** respectively. Station electrical equipment **102i-1** are each illustrated in FIG. 1 as distinct elements for simplified illustrative purposes only and may, alternatively or additionally, be distributed throughout the power generation equipment, **102e-h**, respectively. For example, at wind power station **102c**, each wind turbine may include transformers, frequency converters, power converters, and/or electrical filters. Energy generated at each wind turbine may be collected by distribution lines along strings of wind turbines and move through

2

collectors, switches, transformers, frequency converters, power converters, electrical filters, and/or other station electrical equipment before leaving the wind power station **102c**. Similarly, at photovoltaic power station **102d**, individual photovoltaic panels and/or arrays of photovoltaic panels may include inverters, transformers, frequency converters, power converters, and/or electrical filters. Energy generated at each photovoltaic panel and/or array may be collected by distribution lines along the photovoltaic panels and move through collectors, switches, transformers, frequency converters, power converters, electrical filters, and/or other station electrical equipment before leaving the photovoltaic power station **102d**.

Each generation station **102a-d** may produce AC or DC electrical current which is then typically stepped up to a higher AC voltage before leaving the respective generation station. For example, wind turbines may typically produce AC electrical energy at 600V to 700V, which may then be stepped up to 34.5 kV before leaving the generation station **102d**. In some cases, the voltage may be stepped up multiple times and to a different voltage before exiting the generation station **102c**. As another example, photovoltaic arrays may produce DC voltage at 600V to 900V, which is then inverted to AC voltage and may be stepped up to 34.5 kV before leaving the generation station **102d**. In some cases, the voltage may be stepped up multiple times and to a different voltage before exiting the generation station **102d**.

Upon exiting the generation segment **102**, electrical power generated at generation stations **102a-d** passes through a respective Point of Interconnection ("POI") **103** between a generation station (e.g., **102a-d**) and the rest of the grid. A respective POI **103** represents the point of connection between a generation station's (e.g., **102a-d**) equipment and a transmission system (e.g., transmission segment **104**) associated with electrical grid. In some cases, at the POI **103**, generated power from generation stations **102a-d** may be stepped up at transformer systems **103e-h** to high voltage scales suitable for long-distance transmission along transmission lines **104a**. Typically, the generated electrical energy leaving the POI **103** will be at 115 kV AC or above, but in some cases it may be as low as, for example, 69 kV for shorter distance transmissions along transmission lines **104a**. Each of transformer systems **103e-h** may be a single transformer or may be multiple transformers operating in parallel or series and may be co-located or located in geographically distinct locations. Each of the transformer systems **103e-h** may include substations and other links between the generation stations **102a-d** and the transmission lines **104a**.

A key aspect of the POI **103** is that this is where generation-side metering occurs. One or more utility-scale generation-side meters **103a-d** (e.g., settlement meters) are located at settlement metering points at the respective POI **103** for each generation station **102a-d**. The utility-scale generation-side meters **103a-d** measure power supplied from generation stations **102a-d** into the transmission segment **104** for eventual distribution throughout the grid.

For electricity consumption, the price consumers pay for power distributed through electric power grids is typically composed of, among other costs, Generation, Administration, and Transmission & Distribution ("T&D") costs. T&D costs represent a significant portion of the overall price paid by consumers for electricity. These costs include capital costs (land, equipment, substations, wire, etc.), costs associated with electrical transmission losses, and operation and maintenance costs.

For utility-scale electricity supply, operators of generation stations (e.g., 102a-d) are paid a variable market price for the amount of power the operator generates and provides to the grid, which is typically determined via a power purchase agreement (PPA) between the generation station operator and a grid operator. The amount of power the generation station operator generates and provides to the grid is measured by utility-scale generation-side meters (e.g., 103a-d) at settlement metering points. As illustrated in FIG. 1, the utility-scale generation-side meters 103a-d are shown on a low side of the transformer systems 103e-h, but they may alternatively be located within the transformer systems 103e-h or on the high side of the transformer systems 103e-h. A key aspect of a utility-scale generation-side meter is that it is able to meter the power supplied from a specific generation station into the grid. As a result, the grid operator can use that information to calculate and process payments for power supplied from the generation station to the grid. That price paid for the power supplied from the generation station is then subject to T&D costs, as well as other costs, in order to determine the price paid by consumers.

After passing through the utility-scale generation-side meters in the POI 103, the power originally generated at the generation stations 102a-d is transmitted onto and along the transmission lines 104a in the transmission segment 104. Typically, the electrical energy is transmitted as AC at 115 kV+ or above, though it may be as low as 69 kV for short transmission distances. In some cases, the transmission segment 104 may include further power conversions to aid in efficiency or stability. For example, transmission segment 104 may include high-voltage DC ("HVDC") portions (along with conversion equipment) to aid in frequency synchronization across portions of the transmission segment 104. As another example, transmission segment 104 may include transformers to step AC voltage up and then back down to aid in long distance transmission (e.g., 230 kV, 500 kV, 765 kV, etc.).

Power generated at the generation stations 104a-d is ultimately destined for use by consumers connected to the grid. Once the energy has been transmitted along the transmission segment 104, the voltage will be stepped down by transformer systems 105a-c in the step down segment 105 so that it can move into the distribution segment 106.

In the distribution segment 106, distribution networks 106a-c take power that has been stepped down from the transmission lines 104a and distribute it to local customers, such as local sub-grids (illustrated at 106a), industrial customers, including large EV charging networks (illustrated at 106b), and/or residential and retail customers, including individual EV charging stations (illustrated at 106c). Customer meters 106d, 106f measure the power used by each of the grid-connected customers in distribution networks 106a-c. Customer meters 106d are typically load meters that are unidirectional and measure power use. Some of the local customers in the distribution networks 106a-d may have local wind or solar power systems 106e owned by the customer. As discussed above, these local customer power systems 106e are decentralized and supply power directly to the customer(s). Customers with decentralized wind or solar power systems 106e may have customer meters 106f that are bidirectional or net-metering meters that can track when the local customer power systems 106e produce power in excess of the customer's use, thereby allowing the utility to provide a credit to the customer's monthly electricity bill. Customer meters 106d, 106f differ from utility-scale generation-side meters (e.g., settlement meters) in at least the following characteristics: design (electro-mechanical or electronic vs

current transformer), scale (typically less than 1600 amps vs. typically greater than 50 MW; typically less than 600V vs. typically greater than 14 kV), primary function (use vs. supply metering), economic purpose (credit against use vs. payment for power), and location (in a distribution network at point of use vs. at a settlement metering point at a Point of Interconnection between a generation station and a transmission line).

To maintain stability of the grid, the grid operator strives to maintain a balance between the amount of power entering the grid from generation stations (e.g., 102a-d) and the amount of grid power used by loads (e.g., customers in the distribution segment 106). In order to maintain grid stability and manage congestion, grid operators may take steps to reduce the supply of power arriving from generation stations (e.g., 102a-d) when necessary (e.g., curtailment). Particularly, grid operators may decrease the market price paid for generated power to dis-incentivize generation stations (e.g., 102a-d) from generating and supplying power to the grid. In some cases, the market price may even go negative such that generation station operators must pay for power they allow into the grid. In addition, some situations may arise where grid operators explicitly direct a generation station (e.g., 102a-d) to reduce or stop the amount of power the station is supplying to the grid.

Power market fluctuations, power system conditions (e.g., power factor fluctuation or generation station startup and testing), and operational directives resulting in reduced or discontinued generation all can have disparate effects on renewable energy generators and can occur multiple times in a day and last for indeterminate periods of time. Curtailment, in particular, is particularly problematic.

According to the National Renewable Energy Laboratory's Technical Report TP-6A20-60983 (March 2014):

[C]urtailment [is] a reduction in the output of a generator from what it could otherwise produce given available resources (e.g., wind or sunlight), typically on an involuntary basis. Curtailments can result when operators or utilities command wind and solar generators to reduce output to minimize transmission congestion or otherwise manage the system or achieve the optimal mix of resources. Curtailment of wind and solar resources typically occurs because of transmission congestion or lack of transmission access, but it can also occur for reasons such as excess generation during low load periods that could cause baseload generators to reach minimum generation thresholds, because of voltage or interconnection issues, or to maintain frequency requirements, particularly for small, isolated grids. Curtailment is one among many tools to maintain system energy balance, which can also include grid capacity, hydropower and thermal generation, demand response, storage, and institutional changes. Deciding which method to use is primarily a matter of economics and operational practice.

"Curtailment" today does not necessarily mean what it did in the early 2000s. Two separate changes in the electric sector have shaped curtailment practices since that time: the utility-scale deployment of wind power, which has no fuel cost, and the evolution of wholesale power markets. These simultaneous changes have led to new operational challenges but have also expanded the array of market-based tools for addressing them. Practices vary significantly by region and market design. In places with centrally-organized wholesale power markets and experience with wind power, manual wind energy curtailment processes are increasingly being

US 10,608,433 B1

5

replaced by transparent offer-based market mechanisms that base dispatch on economics. Market protocols that dispatch generation based on economics can also result in renewable energy plants generating less than what they could potentially produce with available wind or sunlight. This is often referred to by grid operators by other terms, such as “downward dispatch.” In places served primarily by vertically integrated utilities, power purchase agreements (PPAs) between the utility and the wind developer increasingly contain financial provisions for curtailment contingencies.

Some reductions in output are determined by how a wind operator values dispatch versus non-dispatch. Other curtailments of wind are determined by the grid operator in response to potential reliability events. Still other curtailments result from overdevelopment of wind power in transmission-constrained areas.

Dispatch below maximum output (curtailment) can be more of an issue for wind and solar generators than it is for fossil generation units because of differences in their cost structures. The economics of wind and solar generation depend on the ability to generate electricity whenever there is sufficient sunlight or wind to power their facilities.

Because wind and solar generators have substantial capital costs but no fuel costs (i.e., minimal variable costs), maximizing output improves their ability to recover capital costs. In contrast, fossil generators have higher variable costs, such as fuel costs. Avoiding these costs can, depending on the economics of a specific generator, to some degree reduce the financial impact of curtailment, especially if the generator’s capital costs are included in a utility’s rate base.

Curtailment may result in available energy being wasted because solar and wind operators have zero variable cost (which may not be true to the same extent for fossil generation units which can simply reduce the amount of fuel that is being used). With wind generation, in particular, it may also take some time for a wind farm to become fully operational following curtailment. As such, until the time that the wind farm is fully operational, the wind farm may not be operating with optimum efficiency and/or may not be able to provide power to the grid.

SUMMARY

In an example, a system includes a set of computing systems. The set of computing systems is configured to perform computational operations using power from a power grid. The system also includes a control system configured to monitor a set of conditions and, while monitoring the set of conditions, receive first power option data based, at least in part, on a power option agreement. The first power option data specify a first minimum power threshold associated with a first time interval. The control system is further configured to provide first control instructions for the set of computing systems based on a combination of at least a portion of the first power option data and at least one condition of the set of conditions responsive to receiving the first power option data. The first control instructions comprises a first power consumption target for the set of computing systems for the first time interval, and the first power consumption target is equal to or greater than the first minimum power threshold associated with the first time interval. The control system is also configured to, while monitoring the set of conditions, receive second power option data based, at least in part, on the power option

6

agreement. The second power option data specify a second minimum power threshold associated with a second time interval. Responsive to receiving the second power option data, the control system is configured to provide second control instructions for the set of computing systems based on a combination of at least a portion of the second power data and at least one condition of the set of conditions. The second control instructions comprises a second power consumption target for the set of computing systems for the second time interval, and wherein the second power consumption target is equal to or greater than the second minimum power threshold associated with the second time interval.

In another example, a method involves monitoring, at a computing system, a set of conditions, and while monitoring the set of conditions, receiving first power option data based, at least in part, on a power option agreement. The first power option data specify a first minimum power threshold associated with a first time interval. The method further involves, responsive to receiving the first power option data, providing first control instructions for a set of computing systems based on a combination of at least a portion of the first power option data and at least one condition of the set of conditions. The first control instructions comprises a first power consumption target for the set of computing systems for the first time interval, and the first power consumption target is equal to or greater than the first minimum power threshold associated with the first time interval. The method further involves, while monitoring the set of conditions, receiving second power option data based, at least in part, on the power option agreement. The second power option data specify a second minimum power threshold associated with a second time interval. The method also involves, responsive to receiving the second power option data, providing second control instructions for the set of computing systems based on a combination of at least a portion of the second power data and at least one condition of the set of conditions. The second control instructions comprises a second power consumption target for the set of computing systems for the second time interval, and the second power consumption target is equal to or greater than the second minimum power threshold associated with the second time interval.

In yet another example, a system is provided. The system includes a set of computing systems, where the set of computing systems is configured to perform computational operations using power from a power grid. The system also includes a control system configured to monitor a set of conditions and receive power option data based, at least in part, on a power option agreement. The power option data specify: (i) a set of minimum power thresholds, and (ii) a set of time intervals, where each minimum power threshold in the set of minimum power thresholds is associated with a time interval in the set of time intervals. The control system is further configured to, responsive to receiving the power option data, determine a performance strategy for the set of computing systems based on a combination of at least a portion of the power option data and at least one condition in the set of conditions. The performance strategy comprises a power consumption target for the set of computing systems for each time interval in the set of time intervals, where each power consumption target is equal to or greater than the minimum power threshold associated with each time interval. The control system is also configured to provide instructions to the set of computing systems to perform one or more computational operations based on the performance strategy.

In a further example, non-transitory computer-readable medium is described that is configured to store instructions,

US 10,608,433 B1

7

that when executed by a computing system, causes the computing system to perform operations consistent with the method steps described above.

Other aspects of the present invention will be apparent from the following description and claims.

BRIEF DESCRIPTION OF THE FIGURES

FIG. 1 shows a typical electrical grid.

FIG. 2 shows a behind-the-meter arrangement with optional grid power, including one or more flexible datacenters, according to one or more example embodiments.

FIG. 3 shows a block diagram of a remote master control system, according to one or more example embodiments.

FIG. 4 a block diagram of a generation station, according to one or more example embodiments.

FIG. 5 shows a block diagram of a flexible datacenter, according to one or more example embodiments.

FIG. 6A shows a structural arrangement of a flexible datacenter, according to one or more example embodiments.

FIG. 6B shows a set of computing systems arranged in a straight configuration, according to one or more example embodiments.

FIG. 7 shows a control distribution system for a flexible datacenter, according to one or more example embodiments.

FIG. 8 shows a control distribution system for a fleet of flexible datacenters, according to one or more example embodiments.

FIG. 9 shows a queue distribution system for a traditional datacenter and a flexible datacenter, according to one or more example embodiments.

FIG. 10A shows a method of dynamic power consumption at a flexible datacenter using behind-the-meter power, according to one or more example embodiments.

FIG. 10B shows a method of dynamic power delivery at a flexible datacenter using behind-the-meter power, according to one or more example embodiments.

FIG. 11 shows a block diagram of a system for implementing power consumption adjustments based on a power option agreement, according to one or more embodiments.

FIG. 12 shows a graph representing power option data based on a power option agreement, according to one or more embodiments.

FIG. 13 shows a method for implementing power consumption adjustments based on a fixed-duration power option agreement, according to one or more embodiments.

FIG. 14 shows a method for implementing power consumption adjustments based on a dynamic power option agreement, according to one or more embodiments.

DETAILED DESCRIPTION

Disclosed examples will now be described more fully hereinafter with reference to the accompanying drawings, in which some, but not all of the disclosed examples are shown. Different examples may be described and should not be construed as limited to the examples set forth herein.

As discussed above, the market price paid to generation stations for supplying power to the grid often fluctuates due to various factors, including the need to maintain grid stability and based on current demand and usage by connected loads in distribution networks. Due to these factors, situations can arise where generation stations are offered substantially lower prices to deter an over-supply of power to the grid. Although these situations typically exist temporarily, generation stations are sometimes forced to either sell power to the grid at the much lower prices or adjust

8

operations to decrease the amount of power generated. Furthermore, some situations may even require generation stations to incur costs in order to offload power to the grid or to shut down generation temporarily.

The volatility in the market price offered for power supplied to the grid can be especially problematic for some types of generation stations. In particular, wind farms and some other types of renewable resource power producers may lack the ability to quickly adjust operations in response to changes in the market price offered for supplying power to the grid. As a result, power generation and management at some generation stations can be inefficient, which can frequently result in power being sold to the grid at low or negative prices. In some situations, a generation station may even opt to halt power generation temporarily to avoid such unfavorable pricing. As such, the time required to halt and to restart the power generation at a generation station can reduce the generation station's ability to take advantage of rising market prices for power supplied to the grid.

Example embodiments provided herein aim to assist generation stations in managing power generation operations and avoid unfavorable power pricing situations like those described above. In particular, example embodiments may involve providing a load that is positioned behind-the-meter ("BTM") and enabling the load to utilize power received behind-the-meter at a generation station in a timely manner. As a general rule of thumb, BTM power is not subject to traditional T&D costs.

For purposes herein, a generation station is considered to be configured for the primary purpose of generating utility-scale power for supply to the electrical grid (e.g., a Wide Area Synchronous Grid or a North American Interconnect).

In one embodiment, equipment located behind-the-meter ("BTM equipment") is equipment that is electrically connected to a generation station's power generation equipment behind (i.e., prior to) the generation station's POI with an electrical grid.

In one embodiment, behind-the-meter power ("BTM power") is electrical power produced by a generation station's power generation equipment and utilized behind (i.e., prior to) the generation station's POI with an electrical grid.

In another embodiment, equipment may be considered behind-the-meter if it is electrically connected to a generation station that is subject to metering by a utility-scale generation-side meter (e.g., settlement meter), and the BTM equipment receives power from the generation station, but the power received by the BTM equipment from the generation station has not passed through the utility-scale generation-side meter. In one embodiment, the utility-scale generation-side meter for the generation station is located at the generation station's POI. In another embodiment, the utility-scale generation-side meter for the generation station is at a location other than the POI for the generation station—for example, a substation between the generation station and the generation station's POI.

In another embodiment, power may be considered behind-the-meter if it is electrical power produced at a generation station that is subject to metering by a utility-scale generation-side meter (e.g., settlement meter), and the BTM power is utilized before being metered at the utility-scale generation-side meter. In one embodiment, the utility-scale generation-side meter for the generation station is located at the generation station's POI. In another embodiment, the utility-scale generation-side meter for the generation station is at a location other than the POI for the generation station—for example, a substation between the generation station and the generation station's POI.

US 10,608,433 B1

9

In another embodiment, equipment may be considered behind-the-meter if it is electrically connected to a generation station that supplies power to a grid, and the BTM equipment receives power from the generation station that is not subject to T&D charges, but power received from the grid that is supplied by the generation station is subject to T&D charges.

In another embodiment, power may be considered behind-the-meter if it is electrical power produced at a generation station that supplies power to a grid, and the BTM power is not subject to T&D charges before being used by electrical equipment, but power received from the grid that is supplied by the generation station is subject to T&D charges.

In another embodiment, equipment may be considered behind-the-meter if the BTM equipment receives power generated from the generation station and that received power is not routed through the electrical grid before being delivered to the BTM equipment.

In another embodiment, power may be considered behind-the-meter if it is electrical power produced at a generation station, and BTM equipment receives that generated power, and that generated power received by the BTM equipment is not routed through the electrical grid before being delivered to the BTM equipment.

For purposes herein, BTM equipment may also be referred to as a behind-the-meter load ("BTM load") when the BTM equipment is actively consuming BTM power.

Beneficially, where BTM power is not subject to traditional T&D costs, a wind farm or other type of generation station can be connected to BTM loads which can allow the generation station to selectively avoid the adverse or less-than optimal cost structure occasionally associated with supplying power to the grid by shunting generated power to the BTM load.

An arrangement that positions and connects a BTM load to a generation station can offer several advantages. In such arrangements, the generation station may selectively choose whether to supply power to the grid or to the BTM load, or both. The operator of a BTM load may pay to utilize BTM power at a cost less than that charged through a consumer meter (e.g., 106d, 1060 located at a distribution network (e.g., 106a-c) receiving power from the grid. The operator of a BTM load may additionally or alternatively charge less than the market rate to consume excess power generated at the generation station during curtailment. As a result, the generation station may direct generated power based on the "best" price that the generation station can receive during a given time frame, and/or the lowest cost the generation station may incur from negative market pricing during curtailment. The "best" price may be the highest price that the generation station may receive for its generated power during a given duration, but can also differ within embodiments and may depend on various factors, such as a prior PPA.

In one example, by having a behind-the-meter option available, a generation station may transition from supplying all generated power to the grid to supplying some or all generated power to one or more BTM loads when the market price paid for power by grid operators drops below a predefined threshold (e.g., the price that the operator of the BTM load is willing to pay the generation station for power). Thus, by having an alternative option for power consumption (i.e., one or more BTM loads), the generation station can selectively utilize the different options to maximize the price received for generated power. In addition, the generation station may also utilize a BTM load to avoid or reduce

10

the economic impact in situations when supplying power to the grid would result in the generation station incurring a net cost.

Providing BTM power to a load can also benefit the BTM load operator. A BTM load may be able to receive and utilize BTM power received from the generation station at a cost that is lower than the cost for power from the grid (e.g., at a customer meter 106d, 1060. This is primarily due to the avoidance (or significant reduction) in T&D costs and the market effects of curtailment. As indicated above, the generation station may be willing to divert generated power to the BTM load rather than supplying the grid due to changing market conditions, or during maintenance periods, or for other non-market conditions. Thus, some situations may arise where the generation station offers power to the BTM load at a price that is substantially lower than the price available on the grid. Furthermore, in some situations, the BTM load may even be able to obtain and utilize BTM power from a generation station at no cost or even at negative pricing since the generation station may rather supply the BTM load with generated power during a given time range instead of paying a higher price for the grid to take the power or modifying operations to decrease power output.

Another example of cost-effective use of BTM power is when the generation station 202 is selling power to the grid at a negative price that is offset by a production tax credit. In certain circumstances, the value of the production tax credit may exceed the price the generation station 202 would have to pay to the grid power to offload generation's station 202 generated power. Advantageously, one or more flexible datacenters 220 may take the generated power behind-the-meter, thereby allowing the generation station 202 to produce and obtain the production tax credit, while selling less power to the grid at the negative price.

Another example of cost-effective behind-the-meter power is when the generation station 202 is selling power to the grid at a negative price because the grid is oversupplied and/or the generation station 202 is instructed to stand down and stop producing altogether. A grid operator may select and direct certain generation stations to go offline and stop supplying power to the grid. Advantageously, one or more flexible datacenters may be used to take power behind-the-meter, thereby allowing the generation station 202 to stop supplying power to the grid, but still stay online and make productive use of the power generated.

Another example of beneficial behind-the-meter power use is when the generation station 202 is producing power that is, with reference to the grid, unstable, out of phase, or at the wrong frequency, or the grid is already unstable, out of phase, or at the wrong frequency. A grid operator may select certain generation stations to go either offline and stop producing power, or to take corrective action with respect to the grid power stability, phase, or frequency. Advantageously, one or more flexible datacenters 220 may be used to selectively consume power behind-the-meter, thereby allowing the generation station 202 to stop providing power to the grid and/or provide corrective feedback to the grid.

Another example of beneficial behind-the-meter power use is that cost-effective behind-the-meter power availability may occur when the generation station 202 is starting up or testing. Individual equipment in the power generation equipment 210 may be routinely offline for installation, maintenance, and/or service and the individual units must be tested prior to coming online as part of overall power generation equipment 210. During such testing or maintenance time, one or more flexible datacenters may be intermittently

11

powered by the one or more units of the power generation equipment 210 that are offline from the overall power generation equipment 210.

Another example of beneficial behind-the-meter power use is that datacenter control systems at the flexible datacenters 220 may quickly ramp up and ramp down power consumption by computing systems in the flexible datacenters 220 based on power availability from the generation station 202. For instance, if the grid requires additional power and signals the demand via a higher local price for power, the generation station 202 can supply the grid with power nearly instantly by having active flexible datacenters 220 quickly ramp down and turn off computing systems (or switch to a stored energy source), thereby reducing an active BTM load.

Another example of beneficial behind-the-meter power use is in new photovoltaic generation stations 202. For example, it is common to design and build new photovoltaic generation stations with a surplus of power capacity to account for degradation in efficiency of the photovoltaic panels over the life of the generation stations. Excess power availability at the generation station can occur when there is excess local power generation and/or low grid demand. In high incident sunlight situations, a photovoltaic generation station 202 may generate more power than the intended capacity of generation station 202. In such situations, a photovoltaic generation station 202 may have to take steps to protect its equipment from damage, which may include taking one or more photovoltaic panels offline or shunting their voltage to dummy loads or the ground. Advantageously, one or more flexible datacenters (e.g., the flexible datacenters 220) may take power behind-the-meter at the Generation Station 202, thereby allowing the generation station 202 to operate the power generation equipment 210 within operating ranges while the flexible datacenters 220 receive BTM power without transmission or distribution costs.

Thus, for at least the reasons described herein, arrangements that involves providing a BTM load as an alternative option for a generation station to direct its generated power to can serve as a mutually beneficial relationship in which both the generation station and the BTM load can economically benefit. The above-noted examples of beneficial use of BTM power are merely exemplary and are not intended to limit the scope of what one of ordinary skill in the art would recognize as benefits to unutilized BTM power capacity, BTM power pricing, or BTM power consumption.

Within example embodiments described herein, various types of utility-scale power producers may operate as generation stations 202 that are capable of supplying power to one or more loads behind-the-meter. For instance, renewable energy sources (e.g., wind, solar, hydroelectric, wave, water current, tidal), fossil fuel power generation sources (coal, natural gas), and other types of power producers (e.g., nuclear power) may be positioned in an arrangement that enables the intermittent supply of generated power behind-the-meter to one or more BTM loads. One of ordinary skill in the art will recognize that the generation station 202 may vary based on an application or design in accordance with one or more example embodiments.

In addition, the particular arrangement (e.g., connections) between the generation station and one or more BTM loads can vary within examples. In one embodiment, a generation station may be positioned in an arrangement wherein the generation station selectively supplies power to the grid and/or to one or more BTM loads. As such, power cost-analysis and other factors (e.g., predicted weather condi-

12

tions, contractual obligations, etc.) may be used by the generation station, a BTM load control system, a remote master control system, or some other system or enterprise, to selectively output power to either the grid or to one or more BTM loads in a manner that maximizes revenue to the generation station. In such an arrangement, the generation station may also be able to supply both the grid and one or more BTM loads simultaneously. In some instances, the arrangement may be configured to allow dynamic manipulation of the percentage of the overall generated power that is supplied to each option at a given time. For example, in some time periods, the generation station may supply no power to the BTM load.

In addition, the type of loads that are positioned behind-the-meter can vary within example embodiments. In general, a load that is behind-the-meter may correspond to any type of load capable of receiving and utilizing power behind-the-meter from a generation station. Some examples of loads include, but are not limited to, datacenters and electric vehicle (EV) charging stations.

Preferred BTM loads are loads that can be subject to intermittent power supply because BTM power may be available intermittently. In some instances, the generation station may generate power intermittently. For example, wind power station 102c and/or photovoltaic power station 102d may only generate power when resource are available or favorable. Additionally or alternatively, BTM power availability at a generation station may only be available intermittently due to power market fluctuations, power system conditions (e.g., power factor fluctuation or generation station startup and testing), and/or operational directives from grid operators or generation station operators.

Some example embodiments of BTM loads described herein involve using one or more computing systems to serve as a BTM load at a generation station. In particular, the computing system or computing systems may receive power behind-the-meter from the generation station to perform various computational operations, such as processing or storing information, performing calculations, mining for cryptocurrencies, supporting blockchain ledgers, and/or executing applications, etc.

Multiple computing systems positioned behind-the-meter may operate as part of a "flexible" datacenter that is configured to operate only intermittently and to receive and utilize BTM power to carry out various computational operations similar to a traditional datacenter. In particular, the flexible datacenter may include computing systems and other components (e.g., support infrastructure, a control system) configured to utilize BTM power from one or more generation stations. The flexible datacenter may be configured to use particular load ramping abilities (e.g., quickly increase or decrease power usage) to effectively operate during intermittent periods of time when power is available from a generation station and supplied to the flexible datacenter behind-the-meter, such as during situations when supplying generated power to the grid is not favorable for the generation station.

In some instances, the amount of power consumed by the computing systems at a flexible datacenter can be ramped up and down quickly, and potentially with high granularity (i.e., the load can be changed in small increments if desired). This may be done based on monitored power system conditions or other information analyses as discussed herein. As recited above, this can enable a generation station to avoid negative power market pricing and to respond quickly to grid direc-

US 10,608,433 B1

13

tives. And by extension, the flexible datacenter may obtain BTM power at a price lower than the cost for power from the grid.

Various types of computing systems can provide granular power ramping. Preferably, the computing systems can perform computational tasks that are immune to, or not substantially hindered by, frequent interruptions or slow-downs in processing as the computing systems ramp down or up. In some embodiments, a control system may be used to activate or de-activate one or more computing systems in an array of computing systems. For example, the control system may provide control instructions to one or more blockchain miners (e.g., a group of blockchain miners), including instructions for powering on or off, adjusting frequency of computing systems performing operations (e.g., adjusting the processing frequency), adjusting the quantity of operations being performed, and when to operate within a low power mode (if available).

Within examples, a control system may correspond to a specialized computing system or may be a computing system within a datacenter serving in the role of the control system. The location of the control system can vary within examples as well. For instance, the control system may be located at a datacenter or physically separate from the datacenter. In some examples, the control system may be part of a network of control systems that manage computational operations, power consumption, and other aspects of a fleet of datacenters. The fleet of datacenters may include one or more traditional datacenters and/or flexible datacenters.

Some embodiments may involve using one or more control systems to direct time-insensitive (e.g., interruptible) computational tasks to computational hardware, such as central processing units (CPUs) and graphics processing units (GPUs), sited behind the meter, while other hardware is sited in front of the meter (i.e., consuming metered grid power via a customer meter (e.g., 106d, 1060) and possibly remote from the behind-the-meter hardware. As such, parallel computing processes, such as Monte Carlo simulations, batch processing of financial transactions, graphics rendering, machine learning, neural network processing, queued operations, and oil and gas field simulation models, are good candidates for such interruptible computational operations.

FIG. 2 shows a behind-the-meter arrangement with optional grid-power, including one or more flexible datacenters, according to one or more example embodiments. Dark arrows illustrate a typical power delivery direction. Consistent with FIG. 1, the arrangement illustrates a generation station 202 in the generation segment 102 of a Wide-Area Synchronous Grid. The generation station 202 supplies utility-scale power (typically >50 MW) via a generation power connection 250 to the Point of Interconnection 103 between the generation station 202 and the rest of the grid. Typically, the power supplied on connection 250 may be at 34.5 kV AC, but it may be higher or lower. Depending on the voltage at connection 250 and the voltage at transmission lines 104a, a transformer system 203 may step up the power supplied from the generation station 202 to high voltage (e.g., 115 kV+AC) for transmission over connection 252 and onto transmission lines 104a of transmission segment 104. Grid power carried on the transmission segment 104 may be from generation station 202 as well as other generation stations (not shown). Also consistent with FIG. 1, grid power is consumed at one or more distribution networks, including example distribution network 206. Grid power may be taken from the transmission lines 104a via connector 254 and stepped down to distribution network

14

voltages (e.g., typically 4 kV to 26 kV AC) and sent into the distribution networks, such as distribution network 206 via distribution line 256. The power on distribution line 256 may be further stepped down (not shown) before entering individual consumer facilities such as a remote master control system 262 and/or traditional datacenters 260 via customer meters 206A, which may correspond to customer meters 106d in FIG. 1, or customer meters 106f in FIG. 1 if the respective consumer facility includes a local customer power system, such as 106e (not shown in FIG. 2).

Consistent with FIG. 1, power entering the grid from generation station 202 is metered by a utility-scale generation-side meter. A utility-scale generation-side meter 253 is shown on the low side of transformer system 203 and an alternative location is shown as 253A on the high side of transformer system 203. Both locations may be considered settlement metering points for the generation station 202 at the POI 103. Alternatively, a utility-scale generation-side meter for the generation station 202 may be located at another location consistent with the descriptions of such meters provided herein.

Generation station 202 includes power generation equipment 210, which may include, as examples, wind turbines and/or photovoltaic panels. Power generation equipment 210 may further include other electrical equipment, including but not limited to switches, busses, collectors, inverters, and power unit transformers (e.g., transformers in wind turbines).

As illustrated in FIG. 2, generation station 202 is configured to connect with BTM equipment which may function as BTM loads. In the illustrated embodiment of FIG. 2, the BTM equipment includes flexible datacenters 220. Various configurations to supply BTM power to flexible datacenters 220 within the arrangement of FIG. 2 are described herein.

In one configuration, generated power may travel from the power generation equipment 210 over one or more connectors 230A, 230B to one or more electrical busses 240A, 240B, respectively. Each of the connectors 230A, 230B may be a switched connector such that power may be routed independently to 240A and/or 240B. For illustrative purposes only, connector 230B is shown with an open switch, and connector 230A is shown with a closed switch, but either or both may be reversed in some embodiments. Aspects of this configuration can be used in various embodiments when BTM power is supplied without significant power conversion to BTM loads.

In various configurations, the busses 240A and 240B may be separated by an open switch 240C or combined into a common bus by a closed switch 240C.

In another configuration, generated power may travel from the power generation equipment 210 to the high side of a local step-down transformer 214. The generated power may then travel from the low side of the local step-down transformer 214 over one or more connectors 232A, 232B to the one or more electrical busses 240A, 240B, respectively. Each of the connectors 232A, 232B may be a switched connector such that power may be routed independently to 240A and/or 240B. For illustrative purposes only, connector 232A is shown with an open switch, and connector 232B is shown with a closed switch, but either or both may be reversed in some embodiments. Aspects of this configuration can be used when it is preferable to connect BTM power to the power generation equipment 210, but the generated power must be stepped down prior to use at the BTM loads.

In another configuration, generated power may travel from the power generation equipment 210 to the low side of a local step-up transformer 212. The generated power may

US 10,608,433 B1

15

then travel from the high side of the local step-up transformer 212 over one or more connectors 234A, 234B to the one or more electrical busses 240A, 240B, respectively. Each of the connectors 234A, 234B may be a switched connector such that power may be routed independently to 240A and/or 240B. For illustrative purposes only, both connectors 234A, 234B are shown with open switches, but either or both may be closed in some embodiments. Aspects of this configuration can be used when it is preferable to connect BTM power to the outbound connector 250 or the high side of the local step-up transformer 212.

In another configuration, generated power may travel from the power generation equipment 210 to the low side of the local step-up transformer 212. The generated power may then travel from the high side of the local step-up transformer 212 to the high side of local step-down transformer 213. The generated power may then travel from the low side of the local step-down transformer 213 over one or more connectors 236A, 236B to the one or more electrical busses 240A, 240B, respectively. Each of the connectors 236A, 236B may be a switched connector such that power may be routed independently to 240A and/or 240B. For illustrative purposes only, both connectors 236A, 236B are shown with open switches, but either or both may be closed in some embodiments. Aspects of this configuration can be used when it is preferable to connect BTM power to the outbound connector 250 or the high side of the local step-up transformer 212, but the power must be stepped down prior to use at the BTM loads.

In one embodiment, power generated at the generation station 202 may be used to power a generation station control system 216 located at the generation station 202, when power is available. The generation station control system 216 may typically control the operation of the generation station 202. Generated power used at the generation station control system 216 may be supplied from bus 240A via connector 216A and/or from bus 240B via connector 216B. Each of the connectors 216A, 216B may be a switched connector such that power may be routed independently to 240A and/or 240B. While the generation station control system 216 can consume BTM power when powered via bus 240A or bus 240B, the BTM power taken by generation station control system 216 is insignificant in terms of rendering an economic benefit. Further, the generation station control system 216 is not configured to operate intermittently, as it generally must remain always on. Further still, the generation station control system 216 does not have the ability to quickly ramp a BTM load up or down.

In another embodiment, grid power may alternatively or additionally be used to power the generation station control system 216. As illustrated here, metered grid power from a distribution network, such as distribution network 206 for simplicity of illustration purposes only, may be used to power generation station control system 216 over connector 216C. Connector 216C may be a switched connector so that metered grid power to the generation station control system 216 can be switched on or off as needed. More commonly, metered grid power would be delivered to the generation station control system 216 via a separate distribution network (not shown), and also over a switched connector. Any such grid power delivered to the generation station control system 216 is metered by a customer meter 206A and subject to T&D costs.

In another embodiment, when power generation equipment 210 is in an idle or off state and not generating power, grid power may backfeed into generation station 202

16

through POI 103 and such grid power may power the generation station control system 216.

In some configurations, an energy storage system 218 may be connected to the generation station 202 via connector 218A, which may be a switched connector. For illustrative purposes only, connector 218A is shown with an open switch but in some embodiments it may be closed. The energy storage system 218 may be connected to bus 240A and/or bus 240B and store energy produced by the power generation equipment 210. The energy storage system may also be isolated from generation station 202 by switch 242A. In times of need, such as when the power generation equipment in an idle or off state and not generating power, the energy storage system may feed power to, for example, the flexible datacenters 220. The energy storage system may also be isolated from the flexible datacenters 220 by switch 242B.

In a preferred embodiment, as illustrated, power generation equipment 210 supplies BTM power via connector 242 to flexible datacenters 220. The BTM power used by the flexible datacenters 220 was generated by the generation station 202 and did not pass through the POI 103 or utility-scale generation-side meter 253, and is not subject to T&D charges. Power received at the flexible datacenters 220 may be received through respective power input connectors 220A. Each of the respective connectors 220A may be a switched connector that can electrically isolate the respective flexible datacenter 220 from the connector 242. Power equipment 220B may be arranged between the flexible datacenters 220 and the connector 242. The power equipment 220B may include, but is not limited to, power conditioners, unit transformers, inverters, and isolation equipment. As illustrated, each flexible datacenter 220 may be served by a respective power equipment 220B. However, in another embodiment, one power equipment 220B may serve multiple flexible datacenter 220.

In one embodiment, flexible datacenters 220 may be considered BTM equipment located behind-the-meter and electrically connected to the power generation equipment 210 behind (i.e., prior to) the generation station's POI 103 with the rest of the electrical grid.

In one embodiment, BTM power produced by the power generation equipment 210 is utilized by the flexible datacenters 220 behind (i.e., prior to) the generation station's POI with an electrical grid.

In another embodiment, flexible datacenters 220 may be considered BTM equipment located behind-the-meter as the flexible datacenters 220 are electrically connected to the generation station 202, and generation station 202 is subject to metering by utility-scale generation-side meter 253 (or 253A, or another utility-scale generation-side meter), and the flexible datacenters 220 receive power from the generation station 202, but the power received by the flexible datacenters 220 from the generation station 202 has not passed through a utility-scale generation-side meter. In this embodiment, the utility-scale generation-side meter 253 (or 253A) for the generation station 202 is located at the generation station's 202 POI 103. In another embodiment, the utility-scale generation-side meter for the generation station 202 is at a location other than the POI for the generation station 202—for example, a substation (not shown) between the generation station 202 and the generation station's POI 103.

In another embodiment, power from the generation station 202 is supplied to the flexible datacenters 220 as BTM power, where power produced at the generation station 202 is subject to metering by utility-scale generation-side meter

253 (or 253A, or another utility-scale generation-side meter), but the BTM power supplied to the flexible datacenters 220 is utilized before being metered at the utility-scale generation-side meter 253 (or 253A, or another utility-scale generation-side meter). In this embodiment, the utility-scale generation-side meter 253 (or 253A) for the generation station 202 is located at the generation station's 202 POI 103. In another embodiment, the utility-scale generation-side meter for the generation station 202 is at a location other than the POI for the generation station 202—for example, a substation (not shown) between the generation station 202 and the generation station's POI 103.

In another embodiment, flexible datacenters 220 may be considered BTM equipment located behind-the-meter as they are electrically connected to the generation station 202 that supplies power to the grid, and the flexible datacenters 220 receive power from the generation station 202 that is not subject to T&D charges, but power otherwise received from the grid that is supplied by the generation station 202 is subject to T&D charges.

In another embodiment, power from the generation station 202 is supplied to the flexible datacenters 220 as BTM power, where electrical power is generated at the generation station 202 that supplies power to a grid, and the generated power is not subject to T&D charges before being used by flexible datacenters 220, but power otherwise received from the connected grid is subject to T&D charges.

In another embodiment, flexible datacenters 220 may be considered BTM equipment located behind-the-meter because they receive power generated from the generation station 202 intended for the grid, and that received power is not routed through the electrical grid before being delivered to the flexible datacenters 220.

In another embodiment, power from the generation station 202 is supplied to the flexible datacenters 220 as BTM power, where electrical power is generated at the generation station 202 for distribution to the grid, and the flexible datacenters 220 receive that power, and that received power is not routed through the electrical grid before being delivered to the flexible datacenters 220.

In another embodiment, metered grid power may alternatively or additionally be used to power one or more of the flexible datacenters 220, or a portion within one or more of the flexible datacenters 220. As illustrated here for simplicity, metered grid power from a distribution network, such as distribution network 206, may be used to power one or more flexible datacenters 220 over connector 256A and/or 256B. Each of connector 256A and/or 256B may be a switched connector so that metered grid power to the flexible datacenters 220 can be switched on or off as needed. More commonly, metered grid power would be delivered to the flexible datacenters 220 via a separate distribution network (not shown), and also over switched connectors. Any such grid power delivered to the flexible datacenters 220 is metered by customer meters 206A and subject to T&D costs. In one embodiment, connector 256B may supply metered grid power to a portion of one or more flexible datacenters 220. For example, connector 256B may supply metered grid power to control and/or communication systems for the flexible datacenters 220 that need constant power and cannot be subject to intermittent BTM power. Connector 242 may supply solely BTM power from the generation station 202 to high power demand computing systems within the flexible datacenters 220, in which case at least a portion of each flexible datacenters 220 so connected is operating as a BTM load. In another embodiment, connector 256A and/or 256B may supply all power used at one or more of the flexible

datacenters 220, in which case each of the flexible datacenters 220 so connected would not be operating as a BTM load.

In another embodiment, when power generation equipment 210 is in an idle or off state and not generating power, grid power may backfeed into generation station 202 through POI 103 and such grid power may power the flexible datacenters 220.

The flexible datacenters 220 are shown in an example arrangement relative to the generation station 202. Particularly, generated power from the generation station 202 may be supplied to the flexible datacenters 220 through a series of connectors and/or busses (e.g., 232B, 240B, 242, 220A). As illustrated, in other embodiments, connectors between the power generation equipment 210 and other components may be switched open or closed, allowing other pathways for power transfer between the power generation equipment 210 and components, including the flexible datacenters 220. Additionally, the connector arrangement shown is illustrative only and other circuit arrangements are contemplated within the scope of supplying BTM power to a BTM load at generation station 202. For example, there may be more or fewer transformers, or one or more of transformers 212, 213, 214 may be transformer systems with multiple steppings and/or may include additional power equipment including but not limited to power conditioners, filters, switches, inverters, and/or AC/DC-DC/AC isolators. As another example, metered grid power connections to flexible datacenters 220 are shown via both 256A and 256B; however, a single connection may connect one or more flexible datacenters 220 (or power equipment 220B) to metered grid power and the one or more flexible datacenters 220 (or power equipment 220B) may include switching apparatus to direct BTM power and/or metered grid power to control systems, communication systems, and/or computing systems as desired.

In some examples, BTM power may arrive at the flexible datacenters 220 in a three-phase AC format. As such, power equipment (e.g., power equipment 220B) at one or more of the flexible datacenters 220 may enable each flexible datacenter 220 to use one or more phases of the power. For instance, the flexible datacenters 220 may utilize power equipment (e.g., power equipment 220B, or alternatively or additionally power equipment that is part of the flexible datacenter 220) to convert BTM power received from the generation station 202 for use at computing systems at each flexible datacenter 220. In other examples, the BTM power may arrive at one or more of the flexible datacenters 220 as DC power. As such, the flexible datacenters 220 may use the DC power to power computing systems. In some such examples, the DC power may be routed through a DC-to-DC converter that is part of power equipment 220B and/or flexible datacenter 220.

In some configurations, a flexible datacenter 220 may be arranged to only have access to power received behind-the-meter from a generation station 202. In the arrangement of FIG. 2, the flexible datacenters 220 may be arranged only with a connection to the generation station 202 and depend solely on power received behind-the-meter from the generation station 202. Alternatively or additionally, the flexible datacenters 220 may receive power from energy storage system 218.

In some configurations, one or more of the flexible datacenters 220 can be arranged to have connections to multiple sources that are capable of supplying power to a flexible datacenter 220. To illustrate a first example, the flexible datacenters 220 are shown connected to connector 242, which can be connected or disconnected via switches to

19

the energy storage system 218 via connector 218A, the generation station 202 via bus 240B, and grid power via metered connector 256A. In one embodiment, the flexible datacenters 220 may selectively use power received behind-the-meter from the generation station 202, stored power supplied by the energy storage system 218, and/or grid power. For instance, flexible datacenters 220 may use power stored in the energy storage system 218 when costs for using power supplied behind-the-meter from the generation station 202 are disadvantageous. By having access to the energy storage system 218 available, the flexible datacenters 220 may use the stored power and allow the generation station 202 to subsequently refill the energy storage system 218 when cost for power behind-the-meter is low. Alternatively, the flexible datacenters 220 may use power from multiple sources simultaneously to power different components (e.g., a first set and a second set of computing systems). Thus, the flexible datacenters 220 may leverage the multiple connections in a manner that can reduce the cost for power used by the computing systems at the flexible datacenters 220. The flexible datacenters 220 control system or the remote master control system 262 may monitor power conditions and other factors to determine whether the flexible datacenters 220 should use power from either the generation station 202, grid power, the energy storage system 218, none of the sources, or a subset of sources during a given time range. Other arrangements are possible as well. For example, the arrangement of FIG. 2 illustrates each flexible datacenter 220 as connected via a single connector 242 to energy storage system 218, generation station 202, and metered grid power via 256A. However, one or more flexible datacenters 220 may have independent switched connections to each energy source, allowing the one or more flexible datacenters 220 to operate from different energy sources than other flexible datacenters 220 at the same time.

The selection of which power source to use at a flexible datacenter (e.g., the flexible datacenters 220) or another type of BTM load can change based on various factors, such as the cost and availability of power from both sources, the type of computing systems using the power at the flexible datacenters 220 (e.g., some systems may require a reliable source of power for a long period), the nature of the computational operations being performed at the flexible datacenters 220 (e.g., a high priority task may require immediate completion regardless of cost), and temperature and weather conditions, among other possible factors. As such, a datacenter control system at the flexible datacenters 220, the remote master control system 262, or another entity (e.g., an operator at the generation station 202) may also influence and/or determine the source of power that the flexible datacenters 220 use at a given time to complete computational operations.

In some example embodiments, the flexible datacenters 220 may use power from the different sources to serve different purposes. For example, the flexible datacenters 220 may use metered power from grid power to power one or more systems at the flexible datacenters 220 that are configured to be always-on (or almost always on), such as a control and/or communication system and/or one or more computing systems (e.g., a set of computing systems performing highly important computational operations). The flexible datacenters 220 may use BTM power to power other components within the flexible datacenters 220, such as one or more computing systems that perform less critical computational operations.

In some examples, one or more flexible datacenters 220 may be deployed at the generation station 202. In other

20

examples, flexible datacenters 220 may be deployed at a location geographically remote from the generation station 202, while still maintaining a BTM power connection to the generation station 202.

In another example arrangement, the generation station 202 may be connected to a first BTM load (e.g., a flexible datacenter 220) and may supply power to additional BTM loads via connections between the first BTM load and the additional BTM loads (e.g., a connection between a flexible datacenter 220 and another flexible datacenter 220).

The arrangement in FIG. 2, and components included therein, are for non-limiting illustration purposes and other arrangements are contemplated in examples. For instance, in another example embodiment, the arrangement of FIG. 2 may include more or fewer components, such as more BTM loads, different connections between power sources and loads, and/or a different number of datacenters. In addition, some examples may involve one or more components within the arrangement of FIG. 2 being combined or further divided.

Within the arrangement of FIG. 2, a control system, such as the remote master control system 262 or another component (e.g., a control system associated with the grid operator, the generation station control system 216, or a datacenter control system associated with a traditional datacenter or one or more flexible datacenters) may use information to efficiently manage various operations of some of the components within the arrangement of FIG. 2. For example, the remote master control system 262 or another component may manage distribution and execution of computational operations at one or more traditional datacenters 260 and/or flexible datacenters 220 via one or more information-processing algorithms. These algorithms may utilize past and current information in real-time to manage operations of the different components. These algorithms may also make some predictions based on past trends and information analysis. In some examples, multiple computing systems may operate as a network to process information.

Information used to make decisions may include economic and/or power-related information, such as monitored power system conditions. Monitored power system conditions may include one or more of excess power generation at a generation station 202, excess power at a generation station 202 that a connected grid cannot receive, power generation at a generation station 202 subject to economic curtailment, power generation at a generation station 202 subject to reliability curtailment, power generation at a generation station 202 subject to power factor correction, low power generation at a generation station 202, start up conditions at a generation station 202, transient power generation conditions at a generation station 202, or testing conditions where there is an economic advantage to using behind-the-meter power generation at a generation station 202. These different monitored power system conditions can be weighted differently during processing and analysis.

In some examples, the information can include the cost for power from available sources (e.g., BTM power at the generation station 202 versus metered grid power) to enable comparisons to be made which power source costs less. In some instances, the information may include historic prices for power to enable the remote master control system 262 or another system to predict potential future prices in similar situations (e.g., the cost of power tends to trend upwards for grid power during warmer weather and peak-use hours). The information may also indicate the availability of power from the various sources (e.g., BTM power at the generation

US 10,608,433 B1

21

station 262, the energy storage system 218 at the generation station 262, and/or metered grid power).

In addition, the information may also include other data, including information associated with operations at components within the arrangement. For instance, the information may include data associated with performance of operations at the flexible datacenters 220 and the traditional datacenters 260, such as the number of computational tasks currently being performed, the types of tasks being performed (e.g., type of computational operation, time-sensitivity, etc.), the number, types, and capabilities of available computing systems, the amount of computational tasks awaiting performance, and the types of computing systems at one or more datacenters, among others. The information may also include data specifying the conditions at one or more datacenters (e.g., whether or not the temperatures are in a desired range, the amount of power available within an energy storage system such as 218), the amount of computational tasks awaiting performance in the queue of one or more of the datacenters, and the identities of the entities associated with the computational operations at one or more of the datacenters. Entities associated with computational operations may be, for example, owners of the datacenters, customers who purchase computational time at the datacenters, or other entities.

The information used by the remote master control system 262 or another component may include data associated with the computational operations to be performed, such as deadlines, priorities (e.g., high vs. low priority tasks), cost to perform based on required computing systems, the optimal computing systems (e.g., CPU vs GPU vs ASIC; processing unit capabilities, speeds, or frequencies, or instructional sets executable by the processing units) for performing each requested computational task, and prices each entity (e.g., company) is willing to pay for computational operations to be performed or otherwise supported via computing systems at a traditional datacenter 260 or a flexible datacenter 220, among others. In addition, the information may also include other data (e.g., weather conditions at locations of datacenters or power sources, any emergencies associated with a datacenter or power source, or the current value of bids associated with an auction for computational tasks).

The information may be updated in-real time and used to make the different operational decisions within the arrangement of FIG. 2. For instance, the information may help a component (e.g., the remote master control system 262 or a control system at a flexible datacenter 220) determine when to ramp up or ramp down power use at a flexible datacenter 220 or when to switch one or more computing systems at a flexible datacenter 220 into a low power mode or to operate at a different frequency, among other operational adjustments. The information can additionally or alternatively help a component within the arrangement of FIG. 2 to determine when to transfer computational operations between computing systems or between datacenters based on various factors. In some instances, the information may also be used to determine when to temporarily stop performing a computational operation or when to perform a computational operation at multiple sites for redundancy or other reasons. The information may further be used to determine when to accept new computational operations from entities or when to temporarily suspend accepting new tasks to be performed due to lack of computing system availability.

The remote master control system 262 represents a computing system that is capable of obtaining, managing, and using the information described above to manage and over-

22

see one or more operations within the arrangement of FIG. 2. As such, the remote master control system 262 may be one or more computing systems configured to process all, or a subset of, the information described above, such as power, environment, computational characterization, and economic factors to assist with the distribution and execution of computing operations among one or more datacenters. For instance, the remote master control system 262 may be configured to obtain and delegate computational operations among one or more datacenters based on a weighted analysis of a variety of factors, including one or more of the cost and availability of power, the types and availability of the computing systems at each datacenter, current and predicted weather conditions at the different locations of flexible datacenters (e.g., flexible datacenters 220) and generation stations (e.g., generation stations 202), levels of power storage available at one or more energy storage systems (e.g., energy storage system 218), and deadlines and other attributes associated with particular computational operations, among other possible factors. As such, the analysis of information performed by the remote master control system 262 may vary within examples. For instance, the remote master control system 262 may use real-time information to determine whether or not to route a computational operation to a particular flexible datacenter (e.g., a flexible datacenter 220) or to transition a computational operation between datacenters (e.g., from traditional datacenter 260 to a flexible datacenter 220).

As shown in FIG. 2, the generation station 202 may be able to supply power to the grid and/or BTM loads such as flexible datacenters 220. With such a configuration, the generation station 202 may selectively provide power to the BTM loads and/or the grid based on economic and power availability considerations. For example, the generation station 202 may supply power to the grid when the price paid for the power exceeds a particular threshold (e.g., the power price offered by operators of the flexible datacenters 220). In some instances, the operator of a flexible datacenter and the operator of a generation station capable of supplying BTM power to the flexible datacenter may utilize a predefined arrangement (e.g., a contract) that specifies a duration and/or price range when the generation station may supply power to the flexible datacenter.

The remote master control system 262 may be capable of directing one or more flexible datacenters 220 to ramp-up or ramp-down to desired power consumption levels, and/or to control cooperative action of multiple flexible datacenters by determining how to power each individual flexible datacenter 220 in accordance with operational directives.

The configuration of the remote master control system 262 can vary within examples as further discussed with respect to FIGS. 2, 3, and 7-9. The remote master control system 262 may operate as a single computing system or may involve a network of computing systems. Preferably, the remote master control system 262 is implemented across one or more servers in a fault-tolerant operating environment that ensures continuous uptime and connectivity by virtue of its distributed nature. Alternatively, although the remote master control system 262 is shown as a physically separate component arrangement for FIG. 2, the remote master control system 262 may be combined with another component in other embodiments. To illustrate an example, the remote master control system 262 may operate as part of a flexible datacenter (e.g., a computing system or a datacenter control system of the flexible datacenter 220), includ-

23

ing sharing components with a flexible datacenter, sharing power with a flexible datacenter, and/or being co-located with a flexible datacenter.

In addition, the remote master control system 262 may communicate with components within the arrangement of FIG. 2 using various communication technologies, including wired and wireless communication technologies. For instance, the remote master control system 262 may use wired (not illustrated) or wireless communication to communicate with datacenter control systems or other computing systems at the flexible datacenters 220 and the traditional datacenters 260. The remote master control system 262 may also communicate with entities inside or outside the arrangement of FIG. 2 and other components within the arrangement of FIG. 2 via wired or wireless communication. For instance, the remote master control system 262 may use wireless communication to obtain computational operations from entities seeking support for the computational operations at one or more datacenters in exchange for payment. The remote master control system 262 may communicate directly with the entities or may obtain the computational operations from the traditional datacenters 260. For instance, an entity may submit jobs (e.g., computational operations) to one or more traditional datacenters 260. The remote master control system 262 may determine that transferring one or more of the computational operations to a flexible datacenter 220 may better support the transferred computational operations. For example, the remote master control system 262 may determine that the transfer may enable the computational operations to be completed quicker and/or at a lower cost. In some examples, the remote master control system 262 may communicate with the entity to obtain approval prior to transferring the one or more computational operations.

The remote master control system 262 may also communicate with grid operators and/or an operator of generation station 202 to help determine power management strategies when distributing computational operations across the various datacenters. In addition, the remote master control system 262 may communicate with other sources, such as weather prediction systems, historical and current power price databases, and auction systems, etc.

In further examples, the remote master control system 262 or another computing system within the arrangement of FIG. 2 may use wired or wireless communication to submit bids within an auction that involves a bidder (e.g., the highest bid) obtaining computational operations or other tasks to be performed. Particularly, the remote master control system 262 may use the information discussed above to develop bids to obtain computing operations for performance at available computing systems at flexible datacenters (e.g., flexible datacenters 220).

In the example arrangement shown in FIG. 2, the flexible datacenters 220 represent example loads that can receive power behind-the-meter from the generation station 202. In such a configuration, the flexible datacenters 220 may obtain and utilize power behind-the-meter from the generation station 202 to perform various computational operations. Performance of a computational operation may involve one or more computing systems providing resources useful in the computational operation. For instance, the flexible datacenters 220 may include one or more computing systems configured to store information, perform calculations and/or parallel processes, perform simulations, mine cryptocurrencies, and execute applications, among other potential tasks. The computing systems can be specialized or generic and can be arranged at each flexible datacenter 220 in a variety

24

of ways (e.g., straight configuration, zig-zag configuration) as further discussed with respect to FIGS. 6A, 6B. Furthermore, although the example arrangement illustrated in FIG. 2 shows configurations where flexible datacenters 220 serve as BTM loads, other types of loads can be used as BTM loads within examples.

The arrangement of FIG. 2 includes the traditional datacenters 260 coupled to metered grid power. The traditional datacenters 260 using metered grid power to provide computational resources to support computational operations. One or more enterprises may assign computational operations to the traditional datacenters 260 with expectations that the datacenters reliably provide resources without interruption (i.e., non-intermittently) to support the computational operations, such as processing abilities, networking, and/or volatile storage. Similarly, one or more enterprises may also request computational operations to be performed by the flexible datacenters 220. The flexible datacenters 220 differ from the traditional datacenters 260 in that the flexible datacenters 220 are arranged and/or configured to be connected to BTM power, are expected to operate intermittently, and are expected to ramp load (and thus computational capability) up or down regularly in response to control directives. In some examples, the flexible datacenters 220 and the traditional datacenters 260 may have similar configurations and may only differ based on the source(s) of power relied upon to power internal computing systems. Preferably, however, the flexible datacenters 220 include particular fast load ramping abilities (e.g., quickly increase or decrease power usage) and are intended and designed to effectively operate during intermittent periods of time.

FIG. 3 shows a block diagram of the remote master control system 300 according to one or more example embodiments. Remote master control system 262 may take the form of remote master control system 300, or may include less than all components in remote master control system 300, different components than in remote master control system 300, and/or more components than in remote master control system 300.

The remote master control system 300 may perform one or more operations described herein and may include a processor 302, a data storage unit 304, a communication interface 306, a user interface 308, an operations and environment analysis module 310, and a queue system 312. In other examples, the remote master control system 300 may include more or fewer components in other possible arrangements.

As shown in FIG. 3, the various components of the remote master control system 300 can be connected via one or more connection mechanisms (e.g., a connection mechanism 314). In this disclosure, the term "connection mechanism" means a mechanism that facilitates communication between two or more devices, systems, components, or other entities. For instance, a connection mechanism can be a simple mechanism, such as a cable, PCB trace, or system bus, or a relatively complex mechanism, such as a packet-based communication network (e.g., LAN, WAN, and/or the Internet). In some instances, a connection mechanism can include a non-tangible medium (e.g., where the connection is wireless).

As part of the arrangement of FIG. 2, the remote master control system 300 (corresponding to remote master control system 262) may perform a variety of operations, such as management and distribution of computational operations among datacenters, monitoring operational, economic, and environment conditions, and power management. For instance, the remote master control system 300 may obtain

US 10,608,433 B1

25

computational operations from one or more enterprises for performance at one or more datacenters. The remote master control system 300 may subsequently use information to distribute and assign the computational operations to one or more datacenters (e.g., the flexible datacenters 220) that have the resources (e.g., particular types of computing systems and available power) available to complete the computational operations. In some examples, the remote master control system 300 may assign all incoming computational operation requests to the queue system 312 and subsequently assign the queued requests to computing systems based on an analysis of current market and power conditions.

Although the remote master control system 300 is shown as a single entity, a network of computing systems may perform the operations of the remote master control system 300 in some examples. For example, the remote master control system 300 may exist in the form of computing systems (e.g., datacenter control systems) distributed across multiple datacenters.

The remote master control system 300 may include one or more processors 302. As such, the processor 302 may represent one or more general-purpose processors (e.g., a microprocessor) and/or one or more special-purpose processors (e.g., a digital signal processor (DSP)). In some examples, the processor 302 may include a combination of processors within examples. The processor 302 may perform operations, including processing data received from the other components within the arrangement of FIG. 2 and data obtained from external sources, including information such as weather forecasting systems, power market price systems, and other types of sources or databases.

The data storage unit 304 may include one or more volatile, non-volatile, removable, and/or non-removable storage components, such as magnetic, optical, or flash storage, and/or can be integrated in whole or in part with the processor 302. As such, the data storage unit 304 may take the form of a non-transitory computer-readable storage medium, having stored thereon program instructions (e.g., compiled or non-compiled program logic and/or machine code) that, when executed by the processor 302, cause the remote master control system 300 to perform one or more acts and/or functions, such as those described in this disclosure. Such program instructions can define and/or be part of a discrete software application. In some instances, the remote master control system 300 can execute program instructions in response to receiving an input, such as from the communication interface 306, the user interface 308, or the operations and environment analysis module 310. The data storage unit 304 may also store other information, such as those types described in this disclosure.

In some examples, the data storage unit 304 may serve as storage for information obtained from one or more external sources. For example, data storage unit 304 may store information obtained from one or more of the traditional datacenters 260, a generation station 202, a system associated with the grid, and flexible datacenters 220. As examples only, data storage 304 may include, in whole or in part, local storage, dedicated server-managed storage, network attached storage, and/or cloud-based storage, and/or combinations thereof.

The communication interface 306 can allow the remote master control system 300 to connect to and/or communicate with another component according to one or more protocols. For instance, the communication interface 306 may be used to obtain information related to current, future, and past prices for power, power availability, current and predicted

26

weather conditions, and information regarding the different datacenters (e.g., current workloads at datacenters, types of computing systems available within datacenters, price to obtain power at each datacenter, levels of power storage available and accessible at each datacenter, etc.). In an example, the communication interface 306 can include a wired interface, such as an Ethernet interface or a high-definition serial-digital-interface (HD-SDI). In another example, the communication interface 406 can include a wireless interface, such as a cellular, satellite, WiMAX, or WI-FI interface. A connection can be a direct connection or an indirect connection, the latter being a connection that passes through and/or traverses one or more components, such as such as a router, switcher, or other network device. Likewise, a wireless transmission can be a direct transmission or an indirect transmission. The communication interface 306 may also utilize other types of wireless communication to enable communication with datacenters positioned at various locations.

The communication interface 306 may enable the remote master control system 300 to communicate with the components of the arrangement of FIG. 2. In addition, the communication interface 306 may also be used to communicate with the various datacenters, power sources, and different enterprises submitting computational operations for the datacenters to support.

The user interface 308 can facilitate interaction between the remote master control system 300 and an administrator or user, if applicable. As such, the user interface 308 can include input components such as a keyboard, a keypad, a mouse, a touch-sensitive panel, a microphone, and/or a camera, and/or output components such as a display device (which, for example, can be combined with a touch-sensitive panel), a sound speaker, and/or a haptic feedback system. More generally, the user interface 308 can include hardware and/or software components that facilitate interaction between remote master control system 300 and the user of the system.

In some examples, the user interface 308 may enable the manual examination and/or manipulation of components within the arrangement of FIG. 2. For instance, an administrator or user may use the user interface 308 to check the status of, or change, one or more computational operations, the performance or power consumption at one or more datacenters, the number of tasks remaining within the queue system 312, and other operations. As such, the user interface 308 may provide remote connectivity to one or more systems within the arrangement of FIG. 2.

The operations and environment analysis module 310 represents a component of the remote master control system 300 associated with obtaining and analyzing information to develop instructions/directives for components within the arrangement of FIG. 2. The information analyzed by the operations and environment analysis module 310 can vary within examples and may include the information described above with respect predicting and/or directing the use of BTM power. For instance, the operations and environment analysis module 310 may obtain and access information related to the current power state of computing systems operating as part of the flexible datacenters 220 and other datacenters that the remote master control system 300 has access to. This information may be used to determine when to adjust power usage or mode of one or more computing systems. In addition, the remote master control system 300 may provide instructions a flexible datacenter 220 to cause a subset of the computing systems to transition into a low power mode to consume less power while still performing

operations at a slower rate. The remote master control system 300 may also use power state information to cause a set of computing systems at a flexible datacenter 220 to operate at a higher power consumption mode. In addition, the remote master control system 300 may transition computing systems into sleep states or power on/off based on information analyzed by the operations and environment analysis module 310.

In some examples, the operations and environment analysis module 310 may use location, weather, activity levels at the flexible datacenters or the generation station, and power cost information to determine control strategies for one or more components in the arrangement of FIG. 2. For instance, the remote master control system 300 may use location information for one or more datacenters to anticipate potential weather conditions that could impact access to power. In addition, the operations and environment analysis module 310 may assist the remote master control system 300 determine whether to transfer computational operations between datacenters based on various economic and power factors.

The queue system 312 represents a queue capable of organizing computational operations to be performed by one or more datacenters. Upon receiving a request to perform a computational operation, the remote master control system 300 may assign the computational operation to the queue until one or more computing systems are available to support the computational operation. The queue system 312 may be used for organizing and transferring computational tasks in real time.

The organizational design of the queue system 312 may vary within examples. In some examples, the queue system 312 may organize indications (e.g., tags, pointers) to sets of computational operations requested by various enterprises. The queue system 312 may operate as a First-In-First-Out (FIFO) data structure. In a FIFO data structure, the first element added to the queue will be the first one to be removed. As such, the queue system 312 may include one or more queues that operate using the FIFO data structure.

In some examples, one or more queues within the queue system 312 may use other designs of queues, including rules to rank or organize queues in a particular manner that can prioritize some sets of computational operations over others. The rules may include one or more of an estimated cost and/or revenue to perform each set of computational operations, an importance assigned to each set of computational operations, and deadlines for initiating or completing each set of computational operations, among others. Examples using a queue system are further described below with respect to FIG. 9.

In some examples, the remote master control system 300 may be configured to monitor one or more auctions to obtain computational operations for datacenters to support. Particularly, the remote master control system 300 may use resource availability and power prices to develop and submit bids to an external or internal auction system for the right to support particular computational operations. As a result, the remote master control system 300 may identify computational operations that could be supported at one or more flexible datacenters 220 at low costs.

FIG. 4 is a block diagram of a generation station 400, according to one or more example embodiments. Generation station 202 may take the form of generation station 400, or may include less than all components in generation station 400, different components than in generation station 400, and/or more components than in generation station 400. The generation station 400 includes power generation equipment

401, a communication interface 408, a behind-the-meter interface 406, a grid interface 404, a user interface 410, a generation station control system 414, and power transformation equipment 402. The power generation equipment 210 may take the form of power generation equipment 401, or may include less than all components in power generation equipment 401, different components than in power generation equipment 401, and/or more components than in power generation equipment 401. Generation station control system 216 may take the form of generation station control system 414, or may include less than all components in generation station control system 414, different components than in generation station control system 414, and/or more components than in generation station control system 414. Some or all of the components generation station 400 may be connected via a communication interface 516. These components are illustrated in FIG. 4 to convey an example configuration for the generation station 400 (corresponding to generation station 202 shown in FIG. 2). In other examples, the generation station 400 may include more or fewer components in other arrangements.

The generation station 400 can correspond to any type of grid-connected utility-scale power producer capable of supplying power to one or more loads. The size, amount of power generated, and other characteristics of the generation station 400 may differ within examples. For instance, the generation station 400 may be a power producer that provides power intermittently. The power generation may depend on monitored power conditions, such as weather at the location of the generation station 400 and other possible conditions. As such, the generation station 400 may be a temporary arrangement, or a permanent facility, configured to supply power. The generation station 400 may supply BTM power to one or more loads and supply metered power to the electrical grid. Particularly, the generation station 400 may supply power to the grid as shown in the arrangement of FIG. 2.

The power generation equipment 401 represents the component or components configured to generate utility-scale power. As such, the power generation equipment 401 may depend on the type of facility that the generation station 400 corresponds to. For instance, the power generation equipment 401 may correspond to electric generators that transform kinetic energy into electricity. The power generation equipment 401 may use electromagnetic induction to generate power. In other examples, the power generation equipment 401 may utilize electrochemistry to transform chemical energy into power. The power generation equipment 401 may use the photovoltaic effect to transform light into electrical energy. In some examples, the power generation equipment 401 may use turbines to generate power. The turbines may be driven by, for example, wind, water, steam or burning gas. Other examples of power production are possible.

The communication interface 408 can enable the generation station 400 to communicate with other components within the arrangement of FIG. 2. As such, the communication interface 408 may operate similarly to the communication interface 306 of the remote master control system 300 and the communication interface 503 of the flexible datacenter 500.

The generation station control system 414 may be one or more computing systems configured to control various aspects of the generation station 400.

The BTM interface 406 is a module configured to enable the power generation equipment 401 to supply BTM power to one or more loads and may include multiple components.

The arrangement of the BTM interface 406 may differ within examples based on various factors, such as the number of flexible datacenters 220 (or 500) coupled to the generation station 400, the proximity of the flexible datacenters 220 (or 500), and the type of generation station 400, among others. In some examples, the BTM interface 406 may be configured to enable power delivery to one or more flexible datacenters positioned near the generation station 400. Alternatively, the BTM interface 406 may also be configured to enable power delivery to one or more flexible datacenters 220 (or 500) positioned remotely from the generation station 400.

The grid interface 404 is a module configured to enable the power generation equipment 401 to supply power to the grid and may include multiple components. As such, the grid interface 404 may couple to one or more transmission lines (e.g., transmission lines 404a shown in FIG. 2) to enable delivery of power to the grid.

The user interface 410 represents an interface that enables administrators and/or other entities to communicate with the generation station 400. As such, the user interface 410 may have a configuration that resembles the configuration of the user interface 308 shown in FIG. 3. An operator may utilize the user interface 410 to control or monitor operations at the generation station 400.

The power transformation equipment 402 represents equipment that can be utilized to enable power delivery from the power generation equipment 401 to the loads and to transmission lines linked to the grid. Example power transformation equipment 402 includes, but is not limited to, transformers, inverters, phase converters, and power conditioners.

FIG. 5 shows a block diagram of a flexible datacenter 500, according to one or more example embodiments. Flexible datacenters 220 may take the form of flexible datacenter 500, or may include less than all components in flexible datacenter 500, different components than in flexible datacenter 500, and/or more components than in flexible datacenter 500. In the example embodiment shown in FIG. 5, the flexible datacenter 500 includes a power input system 502, a communication interface 503, a datacenter control system 504, a power distribution system 506, a climate control system 508, one or more sets of computing systems 512, and a queue system 514. These components are shown connected by a communication bus 528. In other embodiments, the configuration of flexible datacenter 500 can differ, including more or fewer components. In addition, the components within flexible datacenter 500 may be combined or further divided into additional components within other embodiments.

The example configuration shown in FIG. 5 represents one possible configuration for a flexible datacenter. As such, each flexible datacenter may have a different configuration when implemented based on a variety of factors that may influence its design, such as location and temperature that the location, particular uses for the flexible datacenter, source of power supplying computing systems within the flexible datacenter, design influence from an entity (or entities) that implements the flexible datacenter, and space available for the flexible datacenter. Thus, the embodiment of flexible datacenter 220 shown in FIG. 2 represents one possible configuration for a flexible datacenter out of many other possible configurations.

The flexible datacenter 500 may include a design that allows for temporary and/or rapid deployment, setup, and start time for supporting computational operations. For instance, the flexible datacenter 500 may be rapidly

deployed at a location near a source of generation station power (e.g., near a wind farm or solar farm). Rapid deployment may involve positioning the flexible datacenter 500 at a target location and installing and/or configuring one or more racks of computing systems within. The racks may include wheels to enable swift movement of the computing systems. Although the flexible datacenter 500 could theoretically be placed anywhere, transmission losses may be minimized by locating it proximate to BTM power generation.

The physical construction and layout of the flexible datacenter 500 can vary. In some instances, the flexible datacenter 500 may utilize a metal container (e.g., a metal container 602 shown in FIG. 6A). In general, the flexible datacenter 500 may utilize some form of secure weather-proof housing designed to protect interior components from wind, weather, and intrusion. The physical construction and layout of example flexible datacenters are further described with respect to FIGS. 6A-6B.

Within the flexible datacenter 500, various internal components enable the flexible datacenter 500 to utilize power to perform some form of operations. The power input system 502 is a module of the flexible datacenter 500 configured to receive external power and input the power to the different components via assistance from the power distribution system 506. As discussed with respect to FIG. 2, the sources of external power feeding a flexible datacenter can vary in both quantity and type (e.g., the generation stations 202, 400, grid-power, energy storage systems). Power input system 502 includes a BTM power input sub-system 522, and may additionally include other power input sub-systems (e.g., a grid-power input sub-system 524 and/or an energy storage input sub-system 526). In some instances, the quantity of power input sub-systems may depend on the size of the flexible datacenter and the number and/or type of computing systems being powered. In an example embodiment, the flexible datacenter may use grid power as the primary power supply.

In some embodiments, the power input system 502 may include some or all of flexible datacenter Power Equipment 220B. The power input system 502 may be designed to obtain power in different forms (e.g., single phase or three-phase behind-the-meter alternating current ("AC") voltage, and/or direct current ("DC") voltage). As shown, the power input system 502 includes a BTM power input sub-system 522, a grid power input sub-system 524, and an energy input sub-system 526. These sub-systems are included to illustrate example power input sub-systems that the flexible datacenter 500 may utilize, but other examples are possible. In addition, in some instances, these sub-systems may be used simultaneously to supply power to components of the flexible datacenter 500. The sub-systems may also be used based on available power sources.

In some implementations, the BTM power input sub-system 522 may include one or more AC-to-AC step-down transformers used to step down supplied medium-voltage AC to low voltage AC (e.g., 120V to 600V nominal) used to power computing systems 512 and/or other components of flexible datacenter 500. The power input system 502 may also directly receive single-phase low voltage AC from a generation station as BTM power, from grid power, or from a stored energy system such as energy storage system 218. In some implementations, the power input system 502 may provide single-phase AC voltage to the datacenter control system 504 (and/or other components of flexible datacenter 500) independent of power supplied to computing systems 512 to enable the datacenter control system 504 to perform

US 10,608,433 B1

31

management operations for the flexible datacenter 500. For instance, the grid power input sub-system 524 may use grid power to supply power to the datacenter control system 504 to ensure that the datacenter control system 504 can perform control operations and communicate with the remote master control system 300 (or 262) during situations when BTM power is not available. As such, the datacenter control system 504 may utilize power received from the power input system 502 to remain powered to control the operation of flexible datacenter 500, even if the computational operations performed by the computing system 512 are powered intermittently. In some instances, the datacenter control system 504 may switch into a lower power mode to utilize less power while still maintaining the ability to perform some functions.

The power distribution system 506 may distribute incoming power to the various components of the flexible datacenter 500. For instance, the power distribution system 506 may direct power (e.g., single-phase or three-phase AC) to one or more components within flexible datacenter 500. In some embodiments, the power distribution system 506 may include some or all of flexible datacenter Power Equipment 220B.

In some examples, the power input system 502 may provide three phases of three-phase AC voltage to the power distribution system 506. The power distribution system 506 may controllably provide a single phase of AC voltage to each computing system or groups of computing systems 512 disposed within the flexible datacenter 500. The datacenter control system 504 may controllably select which phase of three-phase nominal AC voltage that power distribution system 506 provides to each computing system 512 or groups of computing systems 512. This is one example manner in which the datacenter control system 504 may modulate power delivery (and load at the flexible datacenter 500) by ramping-up flexible datacenter 500 to fully operational status, ramping-down flexible datacenter 500 to offline status (where only datacenter control system 504 remains powered), reducing load by withdrawing power delivery from, or reducing power to, one or more of the computing systems 512 or groups of the computing systems 512, or modulating power factor correction for the generation station 300 (or 202) by controllably adjusting which phases of three-phase nominal AC voltage are used by one or more of the computing systems 512 or groups of the computing systems 512. The datacenter control system 504 may direct power to certain sets of computing systems based on computational operations waiting for computational resources within the queue system 514. In some embodiments, the flexible datacenter 500 may receive BTM DC power to power the computing systems 512.

One of ordinary skill in the art will recognize that a voltage level of three-phase AC voltage may vary based on an application or design and the type or kind of local power generation. As such, a type, kind, or configuration of the operational AC-to-AC step down transformer (not shown) may vary based on the application or design. In addition, the frequency and voltage level of three-phase AC voltage, single-phase AC voltage, and DC voltage may vary based on the application or design in accordance with one or more embodiments.

As discussed above, the datacenter control system 504 may perform operations described herein, such as dynamically modulating power delivery to one or more of the computing systems 512 disposed within flexible datacenter 500. For instance, the datacenter control system 504 may modulate power delivery to one or more of the computing

32

systems 512 based on various factors, such as BTM power availability or an operational directive from a generation station 262 or 300 control system, a remote master control system 262 or 300, or a grid operator. In some examples, the datacenter control system 504 may provide computational operations to sets of computing systems 512 and modulate power delivery based on priorities assigned to the computational operations. For instance, an important computational operation (e.g., based on a deadline for execution and/or price paid by an entity) may be assigned to a particular computing system or set of computing systems 512 that has the capacity, computational abilities to support the computational operation. In addition, the datacenter control system 504 may also prioritize power delivery to the computing system or set of computing systems 512.

In some example, the datacenter control system 504 may further provide directives to one or more computing systems to change operations in some manner. For instance, the datacenter control system 504 may cause one or more computing systems 512 to operate at a lower or higher frequency, change clock cycles, or operate in a different power consumption mode (e.g., a low power mode). These abilities may vary depending on types of computing systems 512 available at the flexible datacenter 500. As a result, the datacenter control system 504 may be configured to analyze the computing systems 512 available either on a periodic basis (e.g., during initial set up of the flexible datacenter 500) or in another manner (e.g., when a new computational operation is assigned to the flexible datacenter 500).

The datacenter control system 504 may also implement directives received from the remote master control system 262 or 300. For instance, the remote master control system 262 or 300 may direct the flexible datacenter 500 to switch into a low power mode. As a result, one or more of the computing systems 512 and other components may switch to the low power mode in response.

The datacenter control system 504 may utilize the communication interface 503 to communicate with the remote master control system 262 or 300, other datacenter control systems of other datacenters, and other entities. As such, the communication interface 503 may include components and operate similar to the communication interface 306 of the remote master control system 300 described with respect to FIG. 4.

The flexible datacenter 500 may also include a climate control system 508 to maintain computing systems 512 within a desired operational temperature range. The climate control system 508 may include various components, such as one or more air intake components, an evaporative cooling system, one or more fans, an immersive cooling system, an air conditioning or refrigerant cooling system, and one or more air outtake components. One of ordinary skill in the art will recognize that any suitable heat extraction system configured to maintain the operation of computing systems 512 within the desired operational temperature range may be used.

The flexible datacenter 500 may further include an energy storage system 510. The energy storage system 510 may store energy for subsequent use by computing systems 512 and other components of flexible datacenter 500. For instance, the energy storage system 510 may include a battery system. The battery system may be configured to convert AC voltage to DC voltage and store power in one or more storage cells. In some instances, the battery system may include a DC-to-AC inverter configured to convert DC

US 10,608,433 B1

33

voltage to AC voltage, and may further include an AC phase-converter, to provide AC voltage for use by flexible datacenter 500.

The energy storage system 510 may be configured to serve as a backup source of power for the flexible datacenter 500. For instance, the energy storage system 510 may receive and retain power from a BTM power source at a low cost (or no cost at all). This low-cost power can then be used by the flexible datacenter 500 at a subsequent point, such as when BTM power costs more. Similarly, the energy storage system 510 may also store energy from other sources (e.g., grid power). As such, the energy storage system 510 may be configured to use one or more of the sub-systems of the power input system 502.

In some examples, the energy storage system 510 may be external to the flexible datacenter 500. For instance, the energy storage system 510 may be an external source that multiple flexible datacenters utilize for back-up power.

The computing systems 512 represent various types of computing systems configured to perform computational operations. Performance of computational operations include a variety of tasks that one or more computing systems may perform, such as data storage, calculations, application processing, parallel processing, data manipulation, cryptocurrency mining, and maintenance of a distributed ledger, among others. As shown in FIG. 5, the computing systems 512 may include one or more CPUs 516, one or more GPUs 518, and/or one or more Application-Specific Integrated Circuits (ASIC's) 520. Each type of computing system 512 may be configured to perform particular operations or types of operations.

Due to different performance features and abilities associated with the different types of computing systems, the datacenter control system 504 may determine, maintain, and/or relay this information about the types and/or abilities of the computing systems, quantity of each type, and availability to the remote master control system 262 or 300 on a routine basis (e.g., periodically or on-demand). This way, the remote master control system 262 or 300 may have current information about the abilities of the computing systems 512 when distributing computational operations for performance at one or more flexible datacenters. Particularly, the remote master control system 262 or 300 may assign computational operations based on various factors, such as the types of computing systems available and the type of computing systems required by each computing operation, the availability of the computing systems, whether computing systems can operate in a low power mode, and/or power consumption and/or costs associated with operating the computing systems, among others.

The quantity and arrangement of these computing systems 512 may vary within examples. In some examples, the configuration and quantity of computing systems 512 may depend on various factors, such as the computational tasks that are performed by the flexible datacenter 500. In other examples, the computing systems 512 may include other types of computing systems as well, such as DSPs, SIMDs, neural processors, and/or quantum processors.

As indicated above, the computing systems 512 can perform various computational operations, including in different configurations. For instance, each computing system may perform a particular computational operation unrelated to the operations performed at other computing systems. Groups of the computing systems 512 may also be used to work together to perform computational operations.

In some examples, multiple computing systems may perform the same computational operation in a redundant

34

configuration. This redundant configuration creates a backup that prevents losing progress on the computational operation in situations of a computing failure or intermittent operation of one or more computing systems. In addition, the computing systems 512 may also perform computational operations using a check point system. The check point system may enable a first computing system to perform operations up to a certain point (e.g., a checkpoint) and switch to a second computing system to continue performing the operations from that certain point. The check point system may also enable the datacenter control system 504 to communicate statuses of computational operations to the remote master control system 262 or 300. This can further enable the remote master control system 262 or 300 to transfer computational operations between different flexible datacenters allowing computing systems at the different flexible datacenters to resume support of computational operations based on the check points.

The queue system 514 may operate similar to the queue system 312 of the remote master control system 300 shown in FIG. 3. Particularly, the queue system 514 may help store and organize computational tasks assigned for performance at the flexible datacenter 500. In some examples, the queue system 514 may be part of a distributed queue system such that each flexible datacenter in a fleet of flexible datacenter includes a queue, and each queue system 514 may be able to communicate with other queue systems. In addition, the remote master control system 262 or 300 may be configured to assign computational tasks to the queues located at each flexible datacenter (e.g., the queue system 514 of the flexible datacenter 500). As such, communication between the remote master control system 262 or 300 and the datacenter control system 504 and/or the queue system 514 may allow organization of computational operations for the flexible datacenter 500 to support.

FIG. 6A shows another structural arrangement for a flexible datacenter, according to one or more example embodiments. The particular structural arrangement shown in FIG. 6A may be implemented at flexible datacenter 500. The illustration depicts the flexible datacenter 500 as a mobile container 702 equipped with the power input system 502, the power distribution system 506, the climate control system 508, the datacenter control system 504, and the computing systems 512 arranged on one or more racks 604. These components of flexible datacenter 500 may be arranged and organized according to an example structural region arrangement. As such, the example illustration represents one possible configuration for the flexible datacenter 500, but others are possible within examples.

As discussed above, the structural arrangement of the flexible datacenter 500 may depend on various factors, such as the ability to maintain temperature within the mobile container 602 within a desired temperature range. The desired temperature range may depend on the geographical location of the mobile container 602 and the type and quantity of the computing systems 512 operating within the flexible datacenter 500 as well as other possible factors. As such, the different design elements of the mobile container 602 including the inner contents and positioning of components may depend on factors that aim to maximize the use of space within mobile container 602, lower the amount of power required to cool the computing systems 512, and make setup of the flexible datacenter 500 efficient. For instance, a first flexible datacenter positioned in a cooler geographic region may include less cooling equipment than a second flexible datacenter positioned in a warmer geographic region.

35

As shown in FIG. 6A, the mobile container 602 may be a storage trailer disposed on permanent or removable wheels and configured for rapid deployment. In other embodiments, the mobile container 602 may be a storage container (not shown) configured for placement on the ground and potentially stacked in a vertical or horizontal manner (not shown). In still other embodiments, the mobile container 602 may be an inflatable container, a floating container, or any other type or kind of container suitable for housing a mobile flexible datacenter. As such, the flexible datacenter 500 may be rapidly deployed on site near a source of unutilized behind-the-meter power generation. And in still other embodiments, the flexible datacenter 500 might not include a mobile container. For example, the flexible datacenter 500 may be situated within a building or another type of stationary environment.

FIG. 6B shows the computing systems 512 in a straight-line configuration for installation within the flexible datacenter 500, according to one or more example embodiments. As indicated above, the flexible datacenter 500 may include a plurality of racks 604, each of which may include one or more computing systems 512 disposed therein. As discussed above, the power input system 502 may provide three phases of AC voltage to the power distribution system 506. In some examples, the power distribution system 506 may controllably provide a single phase of AC voltage to each computing system 512 or group of computing systems 512 disposed within the flexible datacenter 500. As shown in FIG. 6B, for purposes of illustration only, eighteen total racks 604 are divided into a first group of six racks 606, a second group of six racks 608, and a third group of six racks 610, where each rack contains eighteen computing systems 512. The power distribution system (506 of FIG. 5) may, for example, provide a first phase of three-phase AC voltage to the first group of six racks 606, a second phase of three-phase AC voltage to the second group of six racks 608, and a third phase of three-phase AC voltage to the third group of six racks 610. In other embodiments, the quantity of racks and computing systems can vary.

FIG. 7 shows a control distribution system 700 of the flexible datacenter 500 according to one or more example embodiments. The system 700 includes a grid operator 702, a generation station control system 216, a remote master control system 300, and a flexible datacenter 500. As such, the system 700 represents one example configuration for controlling operations of the flexible datacenter 500, but other configurations may include more or fewer components in other arrangements.

The datacenter control system 504 may independently or cooperatively with one or more of the generation station control system 414, the remote master control system 300, and/or the grid operator 702 modulate power at the flexible datacenter 500. During operations, the power delivery to the flexible datacenter 500 may be dynamically adjusted based on conditions or operational directives. The conditions may correspond to economic conditions (e.g., cost for power, aspects of computational operations to be performed), power-related conditions (e.g., availability of the power, the sources offering power), demand response, and/or weather-related conditions, among others.

The generation station control system 414 may be one or more computing systems configured to control various aspects of a generation station (not independently illustrated, e.g., 216 or 400). As such, the generation station control system 414 may communicate with the remote master

36

control system 300 over a networked connection 706 and with the datacenter control system 704 over a networked or other data connection 708.

As discussed with respect to FIGS. 2 and 3, the remote master control system 300 can be one or more computing systems located offsite, but connected via a network connection 710 to the datacenter control system 504. The remote master control system 300 may provide supervisory controls or override control of the flexible datacenter 500 or a fleet of flexible datacenters (not shown).

The grid operator 702 may be one or more computing systems that are configured to control various aspects of the power grid (not independently illustrated) that receives power from the generation station. The grid operator 702 may communicate with the generation station control system 300 over a networked or other data connection 712.

The datacenter control system 504 may monitor BTM power conditions at the generation station and determine when a datacenter ramp-up condition is met. The BTM power availability may include one or more of excess local power generation, excess local power generation that the grid cannot accept, local power generation that is subject to economic curtailment, local power generation that is subject to reliability curtailment, local power generation that is subject to power factor correction, conditions where the cost for power is economically viable (e.g., low cost to obtain power), low priced power, situations where local power generation is prohibitively low, start up situations, transient situations, or testing situations where there is an economic advantage to using locally generated behind-the-meter power generation, specifically power available at little to no cost and with no associated transmission or distribution losses or costs. For example, a datacenter control system may analyze future workload and near term weather conditions at the flexible datacenter.

In some instances, the datacenter ramp-up condition may be met if there is sufficient behind-the-meter power availability and there is no operational directive from the generation station control system 414, the remote master control system 300, or the grid operator 702 to go offline or reduce power. As such, the datacenter control system 504 may enable 714 the power input system 502 to provide power to the power distribution system 506 to power the computing systems 512 or a subset thereof.

The datacenter control system 504 may optionally direct one or more computing systems 512 to perform predetermined computational operations (e.g., distributed computing processes). For example, if the one or more computing systems 512 are configured to perform blockchain hashing operations, the datacenter control system 504 may direct them to perform blockchain hashing operations for a specific blockchain application, such as, for example, Bitcoin, Litecoin, or Ethereum. Alternatively, one or more computing systems 512 may be configured to perform high-throughput computing operations and/or high performance computing operations.

The remote master control system 300 may specify to the datacenter control system 504 what sufficient behind-the-meter power availability constitutes, or the datacenter control system 504 may be programmed with a predetermined preference or criteria on which to make the determination independently. For example, in certain circumstances, sufficient behind-the-meter power availability may be less than that required to fully power the entire flexible datacenter 500. In such circumstances, the datacenter control system 504 may provide power to only a subset of computing systems, or operate the plurality of computing systems in a

US 10,608,433 B1

37

lower power mode, that is within the sufficient, but less than full, range of power that is available. In addition, the computing systems 512 may adjust operational frequency, such as performing more or less processes during a given duration. The computing systems 512 may also adjust internal clocks via over-clocking or under-clocking when performing operations.

While the flexible datacenter 500 is online and operational, a datacenter ramp-down condition may be met when there is insufficient or anticipated to be insufficient, behind-the-meter power availability or there is an operational directive from the generation station control system 414, the remote master control system 300, or the grid operator 702. The datacenter control system 504 may monitor and determine when there is insufficient, or anticipated to be insufficient, behind-the-meter power availability. As noted above, sufficiency may be specified by the remote master control system 300 or the datacenter control system 504 may be programmed with a predetermined preference or criteria on which to make the determination independently.

An operational directive may be based on current dispatch-ability, forward looking forecasts for when behind-the-meter power is, or is expected to be, available, economic considerations, reliability considerations, operational considerations, or the discretion of the generation station control system 414, the remote master control system 300, or the grid operator 702. For example, the generation station control system 414, the remote master control system 300, or the grid operator 702 may issue an operational directive to flexible datacenter 500 to go offline and power down. When the datacenter ramp-down condition is met, the datacenter control system 504 may disable power delivery to the plurality of computing systems (e.g., 512). The datacenter control system 504 may disable 714 the power input system 502 from providing power (e.g., three-phase nominal AC voltage) to the power distribution system 506 to power down the computing systems 512 while the datacenter control system 504 remains powered and is capable of returning service to operating mode at the flexible datacenter 500 when behind-the-meter power becomes available again.

While the flexible datacenter 500 is online and operational, changed conditions or an operational directive may cause the datacenter control system 504 to modulate power consumption by the flexible datacenter 500. The datacenter control system 504 may determine, or the generation station control system 414, the remote master control system 300, or the grid operator 702 may communicate, that a change in local conditions may result in less power generation, availability, or economic feasibility, than would be necessary to fully power the flexible datacenter 500. In such situations, the datacenter control system 504 may take steps to reduce or stop power consumption by the flexible datacenter 500 (other than that required to maintain operation of datacenter control system 504).

Alternatively, the generation station control system 414, the remote master control system 300, or the grid operator 702, may issue an operational directive to reduce power consumption for any reason, the cause of which may be unknown. In response, the datacenter control system 504 may dynamically reduce or withdraw power delivery to one or more computing systems 512 to meet the dictate. The datacenter control system 504 may controllably provide three-phase nominal AC voltage to a smaller subset of computing systems (e.g., 512) to reduce power consumption. The datacenter control system 504 may dynamically reduce the power consumption of one or more computing

38

systems by reducing their operating frequency or forcing them into a lower power mode through a network directive.

Similarly, the flexible datacenter 500 may ramp up power consumption based on various conditions. For instance, the datacenter control system 504 may determine, or the generation control system 414, the remote master control system 300, or the grid operator 702 may communicate, that a change in local conditions may result in greater power generation, availability, or economic feasibility. In such situations, the datacenter control system 504 may take steps to increase power consumption by the flexible datacenter 500.

Alternatively, the generation station control system 414, the remote master control system 300, or the grid operator 702, may issue an operational directive to increase power consumption for any reason, the cause of which may be unknown. In response, the datacenter control system 504 may dynamically increase power delivery to one or more computing systems 512 (or operations at the computing systems 512) to meet the dictate. For instance, one or more computing systems 512 may transition into a higher power mode, which may involve increasing power consumption and/or operation frequency.

One of ordinary skill in the art will recognize that datacenter control system 504 may be configured to have a number of different configurations, such as a number or type or kind of the computing systems 512 that may be powered, and in what operating mode, that correspond to a number of different ranges of sufficient and available behind-the-meter power. As such, the datacenter control system 504 may modulate power delivery over a variety of ranges of sufficient and available unutilized behind-the-meter power availability.

FIG. 8 shows a control distribution system 800 of a fleet of flexible datacenters according to one or more example embodiments. The control distribution system 800 of the flexible datacenter 500 shown and described with respect to FIG. 7 may be extended to a fleet of flexible datacenters as illustrated in FIG. 8. For example, a first generation station (not independently illustrated), such as a wind farm, may include a first plurality of flexible datacenters 802, which may be collocated or distributed across the generation station. A second generation station (not independently illustrated), such as another wind farm or a solar farm, may include a second plurality of flexible datacenters 804, which may be collocated or distributed across the generation station. One of ordinary skill in the art will recognize that the number of flexible datacenters deployed at a given station and the number of stations within the fleet may vary based on an application or design in accordance with one or more example embodiments.

The remote master control system 300 may provide directive to datacenter control systems of the fleet of flexible datacenters in a similar manner to that shown and described with respect to FIG. 7, with the added flexibility to make high level decisions with respect to fleet that may be counterintuitive to a given station. The remote master control system 300 may make decisions regarding the issuance of operational directives to a given generation station based on, for example, the status of each generation station where flexible datacenters are deployed, the workload distributed across fleet, and the expected computational demand required for one or both of the expected workload and predicted power availability. In addition, the remote master control system 300 may shift workloads from the first plurality of flexible datacenters 802 to the second plurality of flexible datacenters 804 for any reason, including, for

example, a loss of BTM power availability at one generation station and the availability of BTM power at another generation station. As such, the remote master control system 300 may communicate with the generation station control systems 806A, 806B to obtain information that can be used to organize and distribute computational operations to the fleets of flexible datacenters 802, 804.

FIG. 9 shows a queue distribution arrangement for a traditional datacenter 902 and a flexible datacenter 500, according to one or more example embodiments. The arrangement of FIG. 9 includes a flexible datacenter 500, a traditional datacenter 902, a queue system 312, a set of communication links 916, 918, 920A, 920B, and the remote master control system 300. The arrangement of FIG. 9 represents an example configuration scheme that can be used to distribute computing operations using a queue system 312 between the traditional datacenter 902 and one or more flexible datacenters. In other examples, the arrangement of FIG. 9 may include more or fewer components in other potential configurations. For instance, the arrangement of FIG. 9 may not include the queue system 312 or may include routes that bypass the queue system 312.

The arrangement of FIG. 9 may enable computational operations requested to be performed by entities (e.g., companies). As such, the arrangement of FIG. 9 may use the queue system 312 to organize incoming computational operations requests to enable efficient distribution to the flexible datacenter 500 and the critical traditional datacenter 902. Particularly, the arrangement of FIG. 9 may use the queue system 312 to organize sets of computational operations thereby increasing the speed of distribution and performance of the different computational operations among datacenters. As a result, the use of the queue system 312 may reduce time to complete operations and reduce costs.

In some examples, one or more components, such as the datacenter control system 504, the remote master control system 300, the queue system 312, or the control system 936, may be configured to identify situations that may arise where using the flexible datacenter 500 can reduce costs or increase productivity of the system, as compared to using the traditional datacenter 902 for computational operations. For example, a component within the arrangement of FIG. 9 may identify when using behind-the-meter power to power the computing systems 512 within the flexible datacenter 500 is at a lower cost compared to using the computing systems 934 within the traditional datacenter 902 that are powered by grid power. Additionally, a component in the arrangement of FIG. 9 may be configured to determine situations when offloading computational operations from the traditional datacenter 902 indirectly (i.e., via the queue system 312) or directly (i.e., bypassing the queue system 312) to the flexible datacenter 500 can increase the performance allotted to the computational operations requested by an entity (e.g., reduce the time required to complete time-sensitive computational operations).

In some examples, the datacenter control system 504 may monitor activity of the computing systems 512 within the flexible datacenter 500 and use the respective activity levels to determine when to obtain computational operations from the queue system 312. For instance, the datacenter control system 504 may analyze various factors prior to requesting or accessing a set of computational operations or an indication of the computational operations for the computing systems 512 to perform. The various factors may include power availability at the flexible datacenter 500 (e.g., either stored or from a BTM source), availability of the computing systems 512 (e.g., percentage of computing systems avail-

able), type of computational operations available, estimated cost to perform the computational operations at the flexible datacenter 500, cost for power, cost for power relative to cost for grid power, and instructions from other components within the system, among others. The datacenter control system 504 may analyze one or more of the factors when determining whether to obtain a new set of computational operations for the computing systems 512 to perform. In such a configuration, the datacenter control system 504 manages the activity of the flexible datacenter 500, including determining when to acquire new sets of computational operations when capacity among the computing systems 512 permit.

In other examples, a component (e.g., the remote master control system 300) within the system may assign or distribute one or more sets of computational operations organized by the queue system 312 to the flexible datacenter 500. For example, the remote master control system 300 may manage the queue system 312, including the distribution of computational operations organized by the queue system 312 to the flexible datacenter 500 and the traditional datacenter 902. The remote master control system 300 may utilize to information described with respect to the Figures above to determine when to assign computational operations to the flexible datacenter 500.

The traditional datacenter 902 may include a power input system 930, a power distribution system 932, a datacenter control system 936, and a set of computing systems 934. The power input system 930 may be configured to receive power from a power grid and distribute the power to the computing systems 934 via the power distribution system 932. The datacenter control system 936 may monitor activity of the computing systems 934 and obtain computational operations to perform from the queue system 312. The datacenter control system 936 may analyze various factors prior to requesting or accessing a set of computational operations or an indication of the computational operations for the computing systems 934 to perform. A component (e.g., the remote master control system 300) within the arrangement of FIG. 9 may assign or distribute one or more sets of computational operations organized by the queue system 312 to the traditional datacenter 902.

The communication link 916 represents one or more links that may serve to connect the flexible datacenter 500, the traditional datacenter 902, and other components within the system (e.g., the remote master control system 300, the queue system 312—connections not shown). In particular, the communication link 916 may enable direct or indirect communication between the flexible datacenter 500 and the traditional datacenter 902. The type of communication link 916 may depend on the locations of the flexible datacenter 500 and the traditional datacenter 902. Within embodiments, different types of communication links can be used, including but not limited to WAN connectivity, cloud-based connectivity, and wired and wireless communication links.

The queue system 312 represents an abstract data type capable of organizing computational operation requests received from entities. As each request for computational operations are received, the queue system 312 may organize the request in some manner for subsequent distribution to a datacenter. Different types of queues can make up the queue system 312 within embodiments. The queue system 312 may be a centralized queue that organizes all requests for computational operations. As a centralized queue, all incoming requests for computational operations may be organized by the centralized queue.

41

In other examples, the queue system 312 may be distributed consisting of multiple queue sub-systems. In the distributed configuration, the queue system 312 may use multiple queue sub-systems to organize different sets of computational operations. Each queue sub-system may be used to organize computational operations based on various factors, such as according to deadlines for completing each set of computational operations, locations of enterprises submitting the computational operations, economic value associated with the completion of computational operations, and quantity of computing resources required for performing each set of computational operations. For instance, a first queue sub-system may organize sets of non-intensive computational operations and a second queue sub-system may organize sets of intensive computational operations. In some examples, the queue system 312 may include queue sub-systems located at each datacenter. This way, each datacenter (e.g., via a datacenter control system) may organize computational operations obtained at the datacenter until computing systems are able to start executing the computational operations. In some examples, the queue system 312 may move computational operations between different computing systems or different datacenters in real-time.

Within the arrangement of FIG. 9, the queue system 312 is shown connected to the remote master control system 300 via the communication link 918. In addition, the queue system 312 is also shown connected to the flexible datacenter via the communication 920A and to the traditional datacenter 902 via the communication link 920B. The communication links 918, 920A, 920B may be similar to the communication link 916 and can be various types of communication links within examples.

The queue system 312 may include a computing system configured to organize and maintain queues within the queue system 312. In another example, one or more other components of the system may maintain and support queues within the queue system 312. For instance, the remote master control system 300 may maintain and support the queue system 312. In other examples, multiple components may maintain and support the queue system 312 in a distributed manner, such as a blockchain configuration.

In some embodiments, the remote master control system 300 may serve as an intermediary that facilitates all communication between flexible datacenter 500 and the traditional datacenter 902. Particularly, the traditional datacenter 902 or the flexible datacenter 500 might need to transmit communications to the remote master control system 300 in order to communicate with the other datacenter. As also shown, the remote master control system 300 may connect to the queue system 312 via the communication link 918. Computational operations may be distributed between the queue system 312 and the remote master control system 300 via the communication link 918. The computational operations may be transferred in real-time and mid-performance from one datacenter to another (e.g., from the traditional datacenter 902 to the flexible datacenter 500). In addition, the remote master control system 300 may manage the queue system 312, including providing resources to support queues within the queue system 312.

As a result, the remote master control system 300 may offload some or all of the computational operations assigned to the traditional datacenter 902 to the flexible datacenter 500. This way, the flexible datacenter 500 can reduce overall computational costs by using the behind-the-meter power to provide computational resources to assist traditional datacenter 902. The remote master control system 300 may use the queue system 312 to temporarily store and organize the

42

offloaded computational operations until a flexible datacenter (e.g., the flexible datacenter 500) is available to perform them. The flexible datacenter 500 consumes behind-the-meter power without transmission or distribution costs, which lowers the costs associated with performing computational operations originally assigned to the traditional datacenter 902. The remote master control system 300 may further communicate with the flexible datacenter 500 via communication link 922 and the traditional datacenter 902 via the communication link 924.

FIG. 10A shows method 1000 of dynamic power consumption at a flexible datacenter using behind-the-meter power according to one or more example embodiments. Other example methods may be used to manipulate the power delivery to one or more flexible datacenters.

In step 1010, the datacenter control system, the remote master control system, or another computing system may monitor behind-the-meter power availability. In some embodiments, monitoring may include receiving information or an operational directive from the generation station control system or the grid operator corresponding to behind-the-meter power availability.

In step 1020, the datacenter control system or the remote master control system 300 may determine when a datacenter ramp-up condition is met. In some embodiments, the datacenter ramp-up condition may be met when there is sufficient behind-the-meter power availability and there is no operational directive from the generation station to go offline or reduce power.

In step 1030, the datacenter control system may enable behind-the-meter power delivery to one or more computing systems. In some instances, the remote master control system may directly enable BTM power delivery to computing systems within the flexible system without instructing the datacenter control system.

In step 1040, once ramped-up, the datacenter control system or the remote master control system may direct one or more computing systems to perform predetermined computational operations. In some embodiments, the predetermined computational operations may include the execution of one or more distributed computing processes, parallel processes, and/or hashing functions, among other types of processes.

While operational, the datacenter control system, the remote master control system, or another computing system may receive an operational directive to modulate power consumption. In some embodiments, the operational directive may be a directive to reduce power consumption. In such embodiments, the datacenter control system or the remote master control system may dynamically reduce power delivery to one or more computing systems or dynamically reduce power consumption of one or more computing systems. In other embodiments, the operational directive may be a directive to provide a power factor correction factor. In such embodiments, the datacenter control system or the remote master control system may dynamically adjust power delivery to one or more computing systems to achieve a desired power factor correction factor. In still other embodiments, the operational directive may be a directive to go offline or power down. In such embodiments, the datacenter control system may disable power delivery to one or more computing systems.

FIG. 10B shows method 1050 of dynamic power delivery to a flexible datacenter using behind-the-meter power according to one or more embodiments. In step 1060, the datacenter control system or the remote master control system may monitor behind-the-meter power availability. In

US 10,608,433 B1

43

certain embodiments, monitoring may include receiving information or an operational directive from the generation station control system or the grid operator corresponding to behind-the-meter power availability.

In step 1070, the datacenter control system or the remote master control system may determine when a datacenter ramp-down condition is met. In certain embodiments, the datacenter ramp-down condition may be met when there is insufficient behind-the-meter power availability or anticipated to be insufficient behind-the-meter power availability or there is an operational directive from the generation station to go offline or reduce power.

In step 1080, the datacenter control system may disable behind-the-meter power delivery to one or more computing systems. In step 1090, once ramped-down, the datacenter control system remains powered and in communication with the remote master control system so that it may dynamically power the flexible datacenter when conditions change.

One of ordinary skill in the art will recognize that a datacenter control system may dynamically modulate power delivery to one or more computing systems of a flexible datacenter based on behind-the-meter power availability or an operational directive. The flexible datacenter may transition between a fully powered down state (while the datacenter control system remains powered), a fully powered up state, and various intermediate states in between. In addition, flexible datacenter may have a blackout state, where all power consumption, including that of the datacenter control system is halted. However, once the flexible datacenter enters the blackout state, it will have to be manually rebooted to restore power to datacenter control system. Generation station conditions or operational directives may cause flexible datacenter to ramp-up, reduce power consumption, change power factor, or ramp-down.

FIG. 11 illustrates a block diagram of a system for implementing control strategies based on a power option agreement, according to one or more embodiments. The system 1100 represents an example arrangement that includes a control system (e.g., the remote master control system 262), a load (e.g., one or more of the datacenters 1102, 1104, and 1106), and a power entity 1140, which may establish and operate in accordance with a power option agreement. Additional arrangements are possible within examples.

In general, a power option agreement is an agreement between a power entity 1140 associated with the delivery of power to a load (e.g., a grid operator, power generation station, or local control station) and the load (e.g., the datacenters 1102-1106). As part of the power option agreement, the load (e.g., load operator, contracting agent for the load, semi-automated control system associated with the load, and/or automated control system associated with the load) provides the power entity 1140 with the right, but not obligation, to reduce the amount of power delivered (e.g., grid power) to the load up to an agreed amount of power during an agreed upon time interval. In order to provide the power entity 1140 with this option, the load needs to be using at least the amount of power subject to the option (e.g., a minimum power threshold). For instance, the load may agree to use at least 1 MW of grid power at all times during a specified 24-hour time interval to provide the power entity 1140 with the option of being able to reduce the amount of power delivered to the load by any amount up to 1 MW at any point during the specified 24-hour time interval. The load may grant the power entity 1140 with this option in exchange for a monetary consideration (e.g., receive power

44

at a reduced price and/or monetary payment if the option is exercised by the power entity).

The power option agreement may be used by the power entity 1140 to reserve the right to reduce the amount of grid power delivered to the load during a set time frame (e.g., the next 24 hours). For instance, the power entity 1140 may exercise a predefined power option to reduce the amount of grid power delivered to the load during a time when the grid power may be better redirected to other loads coupled to the power grid. As such, the power entity 1140 may exercise power option agreements to balance loads coupled to the power grid. In some embodiments, a power option agreement may also specify other parameters, such as costs associated with different levels of power consumption and/or maximum power thresholds for the load to operate according to.

To illustrate an example, a power option agreement may specify that a load (e.g., the datacenters 1102-1106) is required to use at least 10 MW or more at all times during the next 12 hours. Thus, the minimum power threshold according to the power option agreement is 10 MW and this minimum power threshold extends across the time interval of the next 12 hours. In order to comply with the agreement, the load must subsequently operate using 10 MW or more power at all times during the next 12 hours. This way, the load can accommodate a situation where the power entity 1140 exercises the option. Particularly, exercising the option may trigger the load to reduce the amount of power it consumes by an amount up to 10 MW at any point during the 12 hour interval. By establishing this power option agreement, the power entity 1140 can manipulate the amount of power consumed at the load during the next 12 hours by up to 10 MW if power needs to be redirected to another load or a reduction in power consumption is needed for other reasons.

In the example arrangement of the system 1100 shown in FIG. 11, one or more of the datacenters (e.g., the flexible datacenters 1102, 1104, and the traditional datacenter 1106) may operate as the load that is subject to a power option agreement. As the load that is subject to the power option agreement, the datacenters 1102-1106 may execute control instructions in accordance with power target consumption targets that meet or exceed the minimum power thresholds based on the power option agreement.

As shown in FIG. 11, each datacenter 1102-1106 may include a set of computing systems configured to perform computational operations using power from one or more power sources (e.g., BTM power, grid power, and/or grid power subject to a power option agreement). In particular, the flexible datacenter 1102 includes computing systems 1108 arranged into a first set 1114A, a second set 1114B, and a third set 1114C, the flexible datacenter 1104 includes computing systems 1110 arranged into a first set 1116A, a second set 1116B, and a third set 1118B, and the traditional datacenter 1106 includes computing systems 1112 arranged into a first set 1118A, a second set 1118B, and a third set 1118C. Each set of computing systems may include various types of computing systems that can operate in one or more modes.

The different sets of computing systems as well as the multiple datacenters are included in FIG. 11 for illustration purposes. In particular, the variety of computing systems represent different configurations that a load may take while operating in accordance with a power option agreement, and each configuration (as detailed herein) may include ramping up or down power consumption and transferring and performing computational operations between sets of comput-

ing systems and/or datacenters. In other examples, the load that is subject to a power option agreement may take on other configurations (e.g., a single datacenter **1102-1106**, and/or a single set of computing systems).

The remote master control system **262** may serve as a control system that can determine performance strategies and provide control instructions to the load (e.g., one or more of the datacenters **1102-1106**). In particular, the remote master control system **262** can monitor conditions in concert with the minimum power thresholds and time intervals (e.g., power option data) set forth in, and/or derived from, one or more power option agreements to determine performance strategies that can enable the load to meet the expectations of the power option agreement(s) while also efficiently using power to accomplish computational operations. In some instances, the remote master control system **262** may also be subject to the power option agreement and may adjust its own power consumption based on the power option agreement (e.g., ramp up or down power consumption based on the defined minimum power thresholds during time intervals).

To establish a power option agreement, the remote master control system **262** (or another computing system) may communicate with the power entity **1140**. For instance, the remote master control system **262** may provide a request (e.g., a signal and/or a bid) to the power entity **1140** and receive the terms of one or more power option agreements, or power option data related to power option agreements (e.g., data such as minimum power thresholds and time intervals, but not all terms contained within a potential power option agreement) in response. In some examples, the remote master control system **262** may evaluate one or more conditions prior to establishing a power option agreement to ensure that the conditions could enable the load (e.g., the datacenters **1102-1106**) to operate in accordance with the power option agreement. For instance, the remote master control system **262** may check the quantity and deadlines associated with computational operations assigned to specific datacenters prior to establishing specific datacenters as a load subject to a power option agreement. In some cases, multiple power option agreements may be established. For example, each datacenter **1102-1106** may be subject to a different power option agreement, which may result in the remote master control system **262** managing the power consumption at each of the datacenters **1102-1106** differently.

Within the system **1100** shown in FIG. **11**, the power entity **1140** may represent any type of power entity associated with the delivery of power to the load that is subject to a power option agreement. For instance, the power entity **1140** may be a local station control system, a grid operator, or a power generation source. As such, the power entity **1140** may establish power option agreements with the loads via communication with the loads and/or the remote master control system **262**. For example, the power entity **1140** may obtain and accept a bid from a load trying to engage in a power option agreement with the power entity **1140**. The power entity **1140** is shown with a power option module **1142**, which may be used to establish power option agreements (e.g., fixed-duration **1144** and/or dynamic **1146**).

Once a power option agreement is established, the remote master control system **262** may obtain power option data from the power entity **1140** (or another source) that specifies the power and time expectations of the power entity **1140**. As shown in FIG. **11**, the power entity **1140** includes a power option module **1142**, which may be used to provide power option data to the remote master control system **262** and/or

the datacenters **1102-1106**. In particular, the power option data may specify the minimum power threshold or thresholds associated with one or more time intervals for the load to operate at in accordance with based on the power option agreement. The power option data may also specify other constraints that the load should operate in accordance with.

In some examples, the power option data may also include an indication of a monetary penalty that would be imposed upon the load for failure to operate as agreed upon for the power option agreement. In addition, the power option data may also include an indication of a monetary benefit provided to the load operating at power consumption levels that are in accordance with a power option agreement. For instance, monetary benefits could include reduced prices for power, credits for power, and/or monetary payments. In addition, the power option data may include further constraints upon power use, such as one or more maximum power thresholds and corresponding time intervals for the maximum power thresholds.

In some embodiments, the power entity **1140** may correspond to a qualified scheduling entity (QSE). A QSE may submit bids and offers on behalf of resource entities (REs) or load serving entities (LSEs), such as retail electric providers (REPs). QSEs may submit offers to sell and/or bids to buy power (energy) in the Day-Ahead Market (e.g., the next 24 hours) and the Real-Time Market. As such, the remote master control system **262** or another computing system may communicate with one or more QSEs to engage and control one or more loads in accordance with one or more power option agreements.

In some examples, a power option agreement may take the form of a fixed duration power option agreement **1144**. The fixed duration power option agreement **1144** may specify a set of minimum power thresholds and a set of time intervals in advance for an upcoming fixed duration of time covered by the agreement. Each minimum power threshold in the set of minimum power thresholds may be associated with a time interval in the set of time intervals. Examples of such association are provided in FIG. **12**. The fixed duration power option agreement may be established in advanced of the time period covered by the set of time intervals to enable the remote master control system **262** to prepare performance strategies for the load (e.g., the datacenter(s)) associated with the power option agreement. Thus, the remote master control system **262** may evaluate the fixed duration power option and other monitored conditions to determine performance strategies for a set of computing systems (e.g., one or more datacenters) during the different intervals that satisfy the minimum power thresholds.

In other examples, a power option agreement may take the form of a dynamic power option agreement **1146**. For a dynamic power option agreement **1146**, minimum power thresholds may be provided to the remote master control system **262** in real-time (or near real-time). For instance, a dynamic power option agreement may specify that the power entity **1140** may provide adjustments to minimum power thresholds and corresponding time intervals in real-time to the remote master control system **262**. For example, a dynamic power option agreement may provide power option data that specifies a minimum power threshold for immediate adjustments (e.g., for the next hour).

In an embodiment, a dynamic power option agreement **1146** may involve repeat communication between the remote master control system **262** and the power entity **1140**. Particularly, the power entity **1140** may provide signals to the remote master control system **262** that request power consumption adjustments to be initiated at one or more

datacenters by the remote master control system 262 over short time intervals, such as across minutes or seconds. For example, the power entity 1140 may communicate to the remote master control system 262 to ramp power consumption down to a particular level within the next 5 minutes. As a result, the remote master control system 262 may provide instructions to one or more datacenters to ramp down power consumption using a linear ramp over the next 5 minutes to meet the particular level specified by the power entity 1140. The remote master control system 262 may monitor the linear ramp down of power consumption and increase or decrease the rate that the datacenter(s) ramp down power use based on projections and updates received from the power entity 1140. As a result, although the ramp down of power consumption may initially be performed in a linear manner to meet a power target threshold, the remote master control system 262 may adjust the rate of power consumption decrease based on updates from the power entity 1140. For example, 25 percent of the overall power consumption ramp down may occur during a first period (e.g., 4 minutes 30 seconds) of the 5 minutes and the remaining 75 percent of the overall power consumption ramp down may occur during the remaining period of the 5 minutes (e.g., the final 30 seconds). The example percentages are included for illustration purposes and can vary within examples based on various parameters, such as additional communication (e.g., adjustments) provided by the power entity 1140.

In further examples, a power option agreement may operate similarly to both a fixed-duration 1144 and a dynamic power option agreement 1146. Particularly, power option data specifying minimum power thresholds and corresponding time intervals may be provided in advance for the entire fixed-duration of time (e.g., the next 24 hours). Additional power option data may then be subsequently provided enabling the remote master control system 262 to make one or more adjustments to accommodate any changes specified within the additional power option data. For instance, additional power option data may indicate that a power entity exercised its option to deliver less power to the load. As a result, the remote master control system may instruct the load to adjust power consumption based on the power entity reducing the power threshold minimum via exercising the option.

As indicated above, the remote master control system 262 may monitor conditions in addition to the constraints set forth in power option data received from the power entity 1140. Particularly, the remote master control system 262 may monitor and analyze a set of conditions (including the power option data) to determine strategies for assigning, transferring, and otherwise managing computational operations using the one or more datacenters 1102-1106. The determined strategies may enable efficient operation by the datacenters while also ensuring that the datacenters operate at target power consumption levels that meet or exceed the minimum power thresholds set forth within one or more power option agreements.

Example monitored conditions include, but are not limited to, power availability 1120, power prices 1122, computing systems parameters 1124, cryptocurrency prices 1126, computational operation parameters 1128, and weather conditions 1129. Power availability 1120 may include determining power consumption ranges at a set of computing systems and/or at one or more datacenters. In addition, power availability 1120 may also involve determining the source or sources of power available at a datacenter. For instance, the remote master control system 262 may identify the types of power sources (e.g., BTM, grid

power, and/or a battery system) that a datacenter has available. Power prices 1122 may involve an analysis of the different costs associated with powering a set of computing systems. For instance, the remote master control system 262 may determine cost of power from the grid without a power option agreement relative to the cost power from the grid under the power option agreement. In addition, the remote master control system 262 may also compare the cost of grid power relative to the cost of BTM power when available at a datacenter. The power prices 1122 may also involve comparing the cost of using power at different datacenters to determine which datacenter may perform computational operations at a lower cost.

Monitoring computing system parameters 1124 may involve determining parameters related to the computing systems at one or more datacenters. For instance, the remote master control system 262 may monitor various parameters of the computing systems at a datacenter, such as the abilities and availability of various computing systems, the status of the queue used to store computational operations awaiting performance by the computing systems. The remote master control system 262 may determine types and operation modes of the computing systems, including which computing systems could operate in different modes (e.g., a higher power or a lower power mode) and/or at different hash rates and/or frequencies. The remote master control system 262 may also estimate when computing systems may complete current computational operations and/or how many computational operations are assigned to computing systems.

Monitoring cryptocurrency prices 1126 may involve monitoring the current price of one or more cryptocurrencies, the hash rate and/or estimated power consumption associated with mining each cryptocurrency, and other factors associated with the cryptocurrencies. The remote master control system 262 may use data related to monitoring cryptocurrency prices 1126 to determine whether using computing systems to mine a cryptocurrency generates more revenue than the cost of power required for performance of the mining operations.

The remote master control system 262 may monitor parameters related to computational operations (e.g., computational operation parameters 1128). For example, the remote master control system 262 may monitor parameters related to the computational operations requiring performance and currently being performed, such quantity of operations, estimated time to complete, cost to perform each computational operation, deadlines and priorities associated with each computational operation. In addition, the remote master control system 262 may analyze computational operations to determine if a particular type of computing system may perform the computational operation better than other types of computing systems.

Monitoring weather conditions 1129 may include monitoring for any potential power generation disruption due to emergencies or other events, and changes in temperatures or weather conditions at power generators or datacenters that could affect power generation. As such, the operations and environment analysis module (or another component) of the remote master control system 262 may be configured to monitor one or more conditions described above.

The performance strategy determined by the remote master control system 262 based on the monitored conditions and/or power option data can include control instructions for the load (e.g., the datacenters and/or one or more sets of computing systems). For instance, a performance strategy can specify operating parameters, such as operating frequen-

cies, power consumption targets, operating modes, power on/off and/or standby states, and other operation aspects for computing systems at a datacenter.

The performance strategy can also involve aspects related to the assignment, transfer, and performance of computational operations at the computing systems. For instance, the performance strategy may specify computational operations to be performed at the computing systems, an order for completing computational operations based on priorities associated with the computational operations, and an identification of which computing systems should perform which computational operations. In some instances, priorities may depend on revenue associated with completing each computational operation and deadlines for each computational operation.

The monitored conditions may enable efficient distribution and performance of computational operations among computing systems at one or more datacenters (e.g., datacenters 1102-1106) in ways that can reduce costs and/or time to perform computational operations, take advantage of availability and abilities of computing systems at the datacenters 1102-1106, and/or take advantage in changes in the cost for power at the datacenters 1102-1106. In addition, the monitored conditions may also involve consideration of the power option data to ensure that the computing systems consume enough power to meet minimum power thresholds set forth in one or more power option agreements.

The various monitored conditions described above as well as other potential conditions may change dynamically and with great frequency. Thus, to enable efficient distribution and performance of the computational operations at the datacenters, the remote master control system 262 may be configured to monitor changes in the various conditions to assist with the efficient management and operations of the computing systems at each datacenter. For instance, the remote master control system 262 may engage in wired or wireless communication 1130 with datacenter control systems (e.g., datacenter control system 504) at each datacenter as well as other sources (e.g., the power entity 1140) to monitor for changes in the conditions.

The remote master control system 262 may analyze the different conditions in real-time to modulate operating attributes of computing systems at one or more of the datacenters. By using the monitored conditions, the remote master control system 262 may increase revenue, decrease costs, and/or increase performance of computational operations via various modifications, such as transferring computational operations between datacenters or sets of computing systems within a datacenter and adjusting performance at one or more sets of computing systems (e.g., switching to a low power mode).

In some examples, the traditional datacenter 1106 may be the load subject to a power option agreement. As such, the remote master control system 262 may factor the power option agreement when determining whether to perform computational operations using the computing systems 1112 at the traditional datacenter 1106 and/or transfer computational operations to the computing systems 1108, 1110 at the flexible datacenters 1102, 1104. For instance, the monitored conditions may indicate that the price of grid power is substantially higher than BTM power. As a result, the remote master control system 262 may transfer a subset of computational operations from the traditional datacenter 1106 to the flexible datacenters 1102, 1104. The traditional datacenter 1106 may still have some computational operations to perform to ensure that the traditional datacenter 1106 is

using enough power to meet the minimum power threshold or thresholds set forth in the power option agreement.

In some examples, the remote master control system 262 may monitor the grid frequency signal received from the power entity 1140. When the frequency of the grid deviates a threshold amount (e.g., 0.036 Hz above or below 60 Hz), the remote master control system 262 may adjust performance strategies at the load. In some cases, the remote master control system 262 may adjust the power consumption at the load, the number of miners (or computing systems) operating at the load, and/or the frequency or hash rate, among other possible changes. The remote master control system may readjust performance strategies at the load in response to receiving additional power option data from the power entity 1140 (e.g., an indication that the frequency of the grid is back to 60 Hz). In addition, the remote master control system 262 may communicate changes in operations at the load to the power entity 1140. This way, the power entity 1140 may obtain confirmation that the load is adjusting in accordance with a power option agreement.

In some embodiments, a power generation source (e.g., the generation station 400 shown in FIG. 4) may enter into a power option agreement with a grid operator, which may provide the grid operator with the option to reduce the amount of power that the power source generator can deliver to the grid during a defined time interval. For instance, a wind generation farm may enter into the power option agreement with the grid operator. In addition, the remote master control system 262 may also enter into a power option agreement with the power generation source (e.g., the wind farm) to provide a load that can receive excess power from the power generation source when the grid operator exercises the option and lowers the amount of power that the power generation source can deliver to the grid. Thus, rather than reducing the amount of power produced, the power generation source could exercise an option in the agreement with remote master control system 262 and redirect excess power to one or more loads (e.g., a set of computing systems) that could ramp up power consumption in response. In such situations, the remote master control system 262 maybe able to use the excess power from the power generation source (e.g., BTM power) to perform operations at one or more loads at a low cost (or no cost at all). In addition, the power generation source may benefit from the power option agreement by directing excess power to the load instead of temporarily halting power production.

In some examples, a power option agreement may depend on parameters associated balancing grid capacity and demand. For instance, power option agreements may incentivize power consumption ramping during periods of peak grid power use.

FIG. 12 shows a graph representing power option data based on a power option agreement, according to one or more embodiments. The graph 1200 shows power option data arranged according to power 1204 over time 1202. As shown in FIG. 12, time 1202 increases along the X-axis and minimum power thresholds 1204 increase along the Y-axis of the graph 1200. In the example embodiment shown in FIG. 12, the time 1202 increases up to a full day (e.g., 24 hours) in 4 hour increments and the power is shown in MW increasing in intervals of 5 MW. The 24 duration and example minimum power thresholds can differ in other embodiments. Particularly, these values may depend on the terms set forth within the power option agreement.

The graph line 1206 represents sets of minimum power thresholds 1206A, 1206B, 1206C that are specified by

power option data based on the power option agreement. As shown, the graph line **1206** extends the entire 24 hour duration, which indicates that the set of time intervals associated with minimum power thresholds add up to 24 hours. In other examples, the power option agreement may not include a minimum power threshold during a portion of the duration.

The graph line **1206** of the graph **1200** is further used to illustrate power consumption levels that one or more loads (e.g., a set of computing systems) operating according to the power option agreement may utilize during the 24 hour duration. Particularly, the power quantities above the graph line **1206** represents power levels that the load(s) may consume from the power grid during the 24 hour duration that would satisfy the requirements (i.e., the minimum power thresholds **1206A-1206C**) set forth by the power option agreement. In particular, the power quantities above the graph line **1206** include any power quantity that meets or exceeds the minimum power threshold at that time. By extension, the power quantities positioned below the graph line **1206** represents the amount of power that the load could be directed to reduce power consumption by per the power option agreement.

To further illustrate, an initial minimum power threshold **1206A** is shown associated with the time interval starting at hour 0 and extending to hour 8. In particular, the minimum power threshold **1206A** is set at 5 MW during this time interval. Thus, based on the power option data shown in FIG. **12**, the loads must be able to operate at a target power consumption level that is equal to or greater than the 5 MW minimum power threshold **1206A** at all times during the time interval extending from hour 0 to hour 8, in order to be able to satisfy the power option if it is exercised for that time interval. Similarly, the power entity could reduce the power consumed by loads by any amount up to 5 MW at any point during the time interval from hour 0 to hour 8 in accordance with the power option agreement. For instance, the power entity could exercise its option at any point during this time interval to reduce the power consumed by the loads by 3 MW as a way to load balance the power grid. In response to the power entity exercising its option, the load may then operate using 3 MW less power and/or another strategy determined by a control system factoring additional conditions (e.g., the price of grid power, the revenue that could be generated from mining a cryptocurrency, and/or parameters associated with computational operations awaiting performance).

As further shown in the graph **1200** illustrated in FIG. **12**, the next minimum power threshold **1206B** is associated with the following time interval, which starts at hour 8 and extends until hour 16. During this time interval (hour 8 to hour 16), the load(s) may consume 10 MW or more power since the minimum power threshold **1206B** is now set at 10 MW as shown on the Y-axis of the graph **1200**. In light of the power option data, a control system may determine and provide a performance strategy to the load (e.g., a set of computing systems) that includes a power consumption target that meets or exceeds the minimum power threshold **1206B** (i.e., 10 MW). The performance strategy may depend on the power option data as well as other possible conditions, such as the price of grid power, the availability of computing systems, and/or the type of computing operations, etc. In addition, the power entity could exercise its option to reduce the amount of power consumed by the load by 10 MW or less as represented by the power levels under the minimum threshold **1206B** that extend during the time interval of hour 8 to hour 16.

The last minimum power threshold **1206C** is associated with the time interval that starts at hour 16 and extends until hour 24. Similar to the initial minimum power threshold **1206A** associated with the beginning of the graph line **1206**, the last minimum power threshold **1206** is also set at 5 MW. As such, at any point during this interval (hour 16 to hour 24) the loads may consume 5 MW or more to operate in accordance with the power option agreement. As discussed above, by operating at 5 MW or more, the load enables the power consumed from the power grid to be reduced any amount from zero up to 5 MW during this time interval.

When determining the power consumption strategy for a load, a computing system (e.g., the remote master control system **262**) may consider various conditions in addition to the power option data received based on one or more power option agreements. Particularly, the computing system may consider and weigh different conditions in addition to the power option data to determine power consumption targets and/or other control instructions for a load. The conditions may include, but are not limited to, the price of grid power, the price of alternative power sources (e.g., BTM power, stored energy), the revenue associated with mining for one or more cryptocurrencies, parameters related to the computational operations requiring performance (e.g., priorities, deadlines, status of the queue organizing the operations, and/or revenue associated with completing each computational operation), parameters related to the set of computing systems (e.g., types and availabilities of computing systems), and other conditions (e.g., penalties if a minimum power threshold is not met and/or monetary benefits from operating under a power option agreement). By weighing various conditions, the computing system may efficiently manage the set of computing systems, including enabling performance of computational operations cost effectively and/or ensuring at that computing systems operate at target power consumption levels that one or more satisfy power option agreements.

In some examples, the computing system may decrease the amount of power that a set of computing systems consumes from one source and while also increasing the amount of power that the set consumes from another source. For instance, the computing system may determine that the price of power grid power is above a threshold price that makes computational operations relatively expensive to perform using grid power. As a result, the computing system may provide control instructions for the computing systems to consume power grid power that matches a minimum power threshold specified by power option data. This may enable the computing systems to satisfy the power option agreement while also avoiding using pricey grid power beyond the minimum amount required per the power option data. In addition, the computing system may instruct some computing systems to switch to a low power mode or temporarily stop until the price of power from the grid decreases. The computing system may instruct one or more computing systems to operate using power from another source (e.g., BTM power and/or stored energy from a battery system) and/or transfer one or more computational operations to another set of computing systems (e.g., a different datacenter).

When the power option agreement is a fixed duration power option agreement, the computing system may receive an indication of all the minimum power thresholds **1206A-1206C** and an indication of the associated time interval altogether and in advance of the duration associated with the power option agreement. By providing all of the minimum power thresholds **1206A-1206C** and the time intervals in

US 10,608,433 B1

53

advance, the computing system may determine a performance strategy for the load that can extend across the entire duration. Particularly, the computing system may factor the minimum power thresholds and associated time intervals as well as other monitored conditions to determine the performance strategy for the total duration. This can enable the computing system to accept and assign computational operations to computing systems in advance while also using a performance strategy that meets the expectations of a power option agreement.

In some examples, the performance strategy determined by the computing system may include control instructions for the set of computing systems to execute if a power option is exercised. For instance, the performance strategy may specify different power consumption targets for the computing systems that depend on whether a power option is exercised during each time interval.

In some instances, the computing system may modify the performance strategy when one or more conditions change enough to warrant a modification. For instance, the computing system may receive an indication of a change in a minimum power threshold (e.g., a decrease in the minimum power threshold) and determine one or more modifications based on the new minimum power threshold and/or other conditions (e.g., a change in the price of power).

In other examples, the power option agreement may be a dynamic power option agreement. Particularly, the load may be subject to a changing minimum power threshold that can vary during a predefined duration associated with the power option agreement. For example, a dynamic power option agreement may specify that the load is subject to a minimum power threshold that may vary from 0 MW up to 5 MW during the next 24 hours and the particular minimum threshold for each hour may depend on power option data received from the power entity during the prior hour. The dynamic power option agreement may further specify the expected response time from the load. For instance, the power option agreement may indicate that an indication of a new minimum power threshold will be provided an hour prior to the start of the minimum power threshold. The computing system, for example, may receive an indication at hour 7 about the increase in the minimum power threshold **1206B** starting at hour 8. The indication may (or may not) specify the total time interval associated with a new minimum power threshold. For instance, the indication received by the computing system may specify that the 10 MW minimum power threshold **1206B** extends from hour 8 until hour 16. In other instances, the power option data may indicate that the computing system should abide by the new minimum power threshold until receiving further power option data indicating a change to another new minimum power threshold.

In some examples, the power option data may arrive at the computing system in an unknown order from the power entity with expectations of swift power consumption adjustments by the load. As a result, the power option agreement may require fast ramping of the load to meet changes. Ramping may involve ramping up or down power consumption as well as ramping operating techniques (e.g., adjusting frequency or operation mode).

In some embodiments, the type of power option power agreement may depend on the delivery and content of power option data provided to the load (or a control system controlling the load). For instance, a computing system may receive minimum power thresholds set across an entire duration associated with a power option agreement in advance when the power option agreement is a fixed-

54

duration power option agreement. In other instances, the computing system may receive power option data dynamically and adjust operations in real-time (or near real-time). For instance, the computing system may receive a series of power option data that each specifies minimum power threshold changes during the duration set forth in the dynamic power option agreement. To illustrate an example, the computing system may receive power option data during hour 1 that specifies the minimum power threshold for hour 2, power option data during hour 2 that specifies the minimum power threshold for hour 3, and so on across the duration of the dynamic power option agreement.

In some examples, the minimum power threshold for a time interval may be zero during the duration of a power option agreement. As such, the load may use any amount of power from the power grid in accordance with the power option agreement, including no power at all during this time interval. When the price for power is high during this time frame, the load may ramp down power usage to zero MW to avoid paying the high price for power while still being in compliance with the power option agreement.

FIG. 13 illustrates a method for implementing control strategies based on a fixed-duration power option agreement, according to one or more embodiments. The method **1300** serves as an example and may include other steps within other embodiments. A control system (e.g., the remote master control system **262**) may be configured to perform one or more steps of the method **1300**. As such, the control system may take various forms of a computing system, such as a mobile computing device, a wearable computing device, a network of computing systems, etc.

At step **1302**, the method **1300** involves monitoring a set of conditions. For instance, a computing system (e.g., a control system) may monitor various conditions that could impact the performance of operations at one or more loads, including the power consumption targets at the loads. The set of monitored conditions may include a variety of information obtained from one or more external sources, such as one or more datacenters, databases, power generation stations, or types of sources.

Some example conditions include, but are not limited to, the price of grid power, the price and availability of alternative power options (e.g. BTM power, and/or stored energy), parameters of the load (e.g., ramping abilities, type of computing systems, operation modes, etc.), parameters of tasks to be performed using the power at the load (e.g., types, deadlines, priorities, and/or revenue associated with computational operations), availability of other computing systems and their associated costs, and/or revenue associated with mining a cryptocurrency. The computing system may monitor one or more of these conditions as well as others.

At step **1304**, the method **1300** involves receiving power option data based, at least in part, on a power option agreement. As discussed above, the computing system (e.g., a remote master control system) may engage in a power option agreement with a power entity. As a result, the computing system may control a load (e.g., a set of computing systems) in accordance with power thresholds and time intervals received from the power entity based on the power option agreement.

In some examples, the power option data may specify a set of minimum power thresholds and a set of time intervals. Each minimum power threshold in the set of minimum power thresholds may be associated with a time interval in the set of time intervals. To illustrate an example, the power option data may specify a first minimum power threshold

US 10,608,433 B1

55

associated with a first time interval and a second minimum power threshold associated with a second time interval, with the second time interval subsequent to the first time interval.

The set of time intervals may add up to the duration represented by the power option agreement. For instance, the total duration of the set of time intervals may correspond to a twenty-four hour period (e.g., the next day). In other examples, the power option agreement may span across a different duration (e.g., 12 hours). In additional embodiments, the power option data may specify other information, such as monetary incentives associated with parameters of the power option agreement and/or one or more maximum power thresholds.

At step **1306**, the method **1300** involves determining a performance strategy for the set of computing systems based on a combination of at least a portion of the power option data and at least one condition in the set of conditions. The performance strategy may be determined responsive to receiving the power option data. In addition, the performance strategy may include a power consumption target for the set of computing systems for each time interval in the set of time intervals. In some examples, each power consumption target is equal to or greater than the minimum power threshold associated with each time interval.

As an example, the performance strategy may specify a first power consumption target for the set of computing systems for a first time interval such that the first power consumption target is equal to or greater than a first minimum power threshold associated with the first time interval and a second power consumption target for the set for a second time interval in a similar manner (i.e., the second power consumption target is equal to or greater than a second minimum power threshold).

In some examples, the performance strategy may include an sequence for the set of computing systems to follow when performing computational operations. The sequence, for example, may be based on priorities associated with the computational operations. In addition, the performance strategy may include one or more power consumption targets that are greater than the minimum power thresholds when the price of power from the power grid is below a threshold price during the time intervals associated with the minimum power thresholds.

The performance strategy may also involve transferring, delaying, or adjusting one or more computational operations performed at the set of computing systems. In addition, the performance strategy may involve adjusting operations at the computing systems. For instance, one or more computing systems may switch modes (e.g., operate at a higher frequency or switch to a low power mode).

In addition, the performance strategy may also specify power consumption targets for the set of computing systems to use if the power option is exercised during an interval. This way, the computing systems may continue to perform computational operations (or suspend performance) based on the power option being exercised.

At step **1308**, the method **1300** involves providing instructions to the set of computing systems to perform one or more computational operations based on the performance strategy. For example, the set of computing systems may operate according to the performance strategy to ensure that the minimum power thresholds are met during the defined time intervals based on the power option agreement.

Some examples may further involve receiving subsequent power option data based, at least in part, on the power option agreement. The subsequent power option data may specify to decrease one or more minimum power thresholds of the

56

set of power thresholds. Responsive to receiving the subsequent power option data, the performance strategy for the set of computing systems may be modified based on a combination of at least a portion of the subsequent power option data and one or more conditions of the monitored conditions. The modified performance strategy may include one or more reduced power consumption targets for the set of computing systems. The amount of the reduction in a power consumption target may depend linearly with the amount that the corresponding minimum power threshold was reduced by. For instance, when a minimum power threshold for a time interval is reduced from 10 MW to 5 MW, the power consumption target for that time interval may be reduced from 10 MW to 5 MW. Instructions may be provided to the set of computing systems to perform computational operations based on the modified performance strategy.

FIG. 14 illustrates a method for implementing control strategies based on a dynamic power option agreement, according to one or more embodiments. The method **1400** serves as an example and may include other steps within other embodiments. Similar to the method **1400**, a control system (e.g., the remote master control system **262**) may be configured to perform one or more steps of the method **1400**. As such, the control system may take various forms of a computing system, such as a mobile computing device, a wearable computing device, a network of computing systems, etc.

At block **1402**, the method **1400** involves monitoring a set of conditions. Similar to block **1302** of the method **1300**, a computing system may monitor various conditions to determine instructions for controlling a set of computing systems.

At block **1404**, the method **1400** involves receiving first power option data based, at least in part, on a power option agreement while monitoring the set of conditions. The first power option data may specify a first minimum power threshold associated with a first time interval. For example, the first power option data may specify a minimum power threshold of 10 MW for the next hour, which may start in an hour or less.

The power option agreement may correspond to a dynamic power option agreement in some examples. When managing a load with respect to a dynamic power option agreement, a computing system may receive power option data specifying changes in minimum power thresholds that a load (e.g., the set of computing systems) may be designated to use in the near term (e.g., the next hour). For example, the computing system may receive power option data during each hour of the duration specified by a power option agreement that indicates a minimum power threshold for the next hour.

At block **1406**, the method **1400** involves providing first control instructions for a set of computing systems based on a combination of at least a portion of the first power option data and at least one condition. The first control instructions may be provided responsive to receiving the first power option data.

The first control instructions may include a first power consumption target for the set of computing systems for the first time interval. Particularly, the first power consumption target may be equal to or greater than the first minimum power threshold associated with the first time interval. For example, the first power consumption target may be greater than the first minimum power threshold when a cost of power from the power grid is below a threshold price during the first time interval. In other instances, the first power consumption target may be equal to the first minimum power

US 10,608,433 B1

57

threshold when the cost of power from the power grid is greater than the threshold price.

In some examples, control instructions may specify a sequence for the computing systems to follow when performing computational operations. The sequence may be based on priorities associated with each computational operation.

The first control instructions may be determined based on a combination of the first power option data, the price of power from the power grid, and parameters associated with computational operations to be performed at the set of computing systems.

In some examples, the first control instructions may involve ramping up or down power consumption at the set of computing systems. The power consumption may be ramped up or down based on the first minimum power threshold and one or more other conditions (e.g., the price of power).

At block 1408, the method 1400 involves receiving second power option data based, at least in part, on the power option agreement while monitoring the set of conditions. The computing system may receive the second power option data subsequent to receiving the first power option data. The second power option data may specify a second minimum power threshold associated with a second time interval. For example, the second minimum power threshold may be 7 MW over the duration of the upcoming hour. In other examples, the second minimum power threshold may differ as shown in FIG. 12.

In some instances, the computing system may receive the second power option data during the first time interval such that the second time interval overlaps the first time interval. For instance, the computing system may receive the second power option data to enable real-time adjustments to be made to the power consumed at the set of computing systems.

At block 1410, the method 1400 involves providing second control instructions for the set of computing systems based on a combination of at least a portion of the second power option data and at least one condition. The second control instructions may be provided responsive to receiving the second power option data. The second control instructions may specify a second power consumption target for the set of computing systems for the second time interval. The second power consumption target may be equal to or greater than the second minimum power threshold associated with the second time interval.

In some examples, the computing system may provide a request to a QSE to determine the power option agreement. As such, the computing system may receive power option data (e.g., the first and second power option data) in response to providing the request to the QSE.

The computing system may monitor the price of power from the power grid, and the global mining hash rate and a price for a cryptocurrency (e.g., Bitcoin), among other conditions. The computing system may determine control instructions (e.g., the first and/or second control instructions) based on a combination of power option data, the price of power from the power grid, and the global mining hash rate and the price for the cryptocurrency. For instance, the computing system may cause one or more computing systems (e.g., a subset of computing systems) to perform mining operations for the cryptocurrency when the price of power from the power grid is equal to or less than a revenue obtained by performing the mining operations for the cryptocurrency.

58

Advantages of one or more embodiments of the present invention may include one or more of the following:

One or more embodiments of the present invention provides a green solution to two prominent problems: the exponential increase in power required for growing blockchain operations and the unutilized and typically wasted energy generated from renewable energy sources.

One or more embodiments of the present invention allows for the rapid deployment of mobile datacenters to local stations. The mobile datacenters may be deployed on site, near the source of power generation, and receive low cost or unutilized power behind-the-meter when it is available.

One or more embodiments of the present invention provide the use of a queue system to organize computational operations and enable efficient distribution of the computational operations across multiple datacenters.

One or more embodiments of the present invention enable datacenters to access and obtain computational operations organized by a queue system.

One or more embodiments of the present invention allows for the power delivery to the datacenter to be modulated based on conditions or an operational directive received from the local station or the grid operator.

One or more embodiments of the present invention may dynamically adjust power consumption by ramping-up, ramping-down, or adjusting the power consumption of one or more computing systems within the flexible datacenter.

One or more embodiments of the present invention may be powered by behind-the-meter power that is free from transmission and distribution costs. As such, the flexible datacenter may perform computational operations, such as distributed computing processes, with little to no energy cost.

One or more embodiments of the present invention provides a number of benefits to the hosting local station. The local station may use the flexible datacenter to adjust a load, provide a power factor correction, to offload power, or operate in a manner that invokes a production tax credit and/or generates incremental revenue.

One or more embodiments of the present invention allows for continued shunting of behind-the-meter power into a storage solution when a flexible datacenter cannot fully utilize excess generated behind-the-meter power.

One or more embodiments of the present invention allows for continued use of stored behind-the-meter power when a flexible datacenter can be operational but there is not an excess of generated behind-the-meter power.

One or more embodiments of the present invention allows for management and distribution of computational operations at computing systems across a fleet of datacenters such that the performance of the computational operations take advantages of increased efficiency and decreased costs.

It will also be recognized by the skilled worker that, in addition to improved efficiencies in controlling power delivery from intermittent generation sources, such as wind farms and solar panel arrays, to regulated power grids, the invention provides more economically efficient control and stability of such power grids in the implementation of the technical features as set forth herein.

While the present invention has been described with respect to the above-noted embodiments, those skilled in the art, having the benefit of this disclosure, will recognize that other embodiments may be devised that are within the scope of the invention as disclosed herein. Accordingly, the scope of the invention should be limited only by the appended claims.

US 10,608,433 B1

59

What is claimed is:

1. A system comprising:
 - a set of computing systems, wherein the set of computing systems is configured to perform computational operations using power from a power grid;
 - a control system configured to:
 - monitor a set of conditions;
 - receive power option data based, at least in part, on a power option agreement, wherein the power option data specify: (i) a set of minimum power thresholds, and (ii) a set of time intervals, wherein each minimum power threshold in the set of minimum power thresholds is associated with a time interval in the set of time intervals;
 - responsive to receiving the power option data, determine a performance strategy for the set of computing systems based on a combination of at least a portion of the power option data and at least one condition in the set of conditions, wherein the performance strategy comprises a power consumption target for the set of computing systems for each time interval in the set of time intervals, wherein each power consumption target is equal to or greater than the minimum power threshold associated with each time interval; and
 - provide instructions to the set of computing systems to perform one or more computational operations based on the performance strategy.
2. The system of claim 1, wherein the control system is configured to monitor the set of conditions comprising:
 - a price of power from the power grid; and
 - a plurality of parameters associated with one or more computational operations to be performed at the set of computing systems.
3. The system of claim 2, wherein the control system is configured to:
 - determine the performance strategy for the set of computing systems based on a combination of at least the portion option data, the price of power from the power grid, and the plurality of parameters associated with the one or more computational operations.
4. The system of claim 3, wherein the performance strategy further comprises:
 - an order for the set of computing systems to follow when performing the one or more computational operations, wherein the order is based on respective priorities associated with the one or more computational operations.
5. The system of claim 4, wherein the performance strategy further comprises:
 - at least one power consumption target that is greater than a minimum power threshold when the price of power from the power grid is below a threshold price during the time interval associated with the minimum power threshold.
6. The system of claim 1, wherein the control system is further configured to:
 - receive subsequent power option data based, at least in part, on the power option agreement,
 - wherein the subsequent power option data specify to decrease one or more minimum power thresholds of the set of minimum power thresholds.
7. The system of claim 6, wherein the control system is further configured to:
 - responsive to receiving the subsequent power option data, modify the performance strategy for the set of computing systems based on a combination of at least the

60

- portion of the subsequent power option data and at least one condition in the set of conditions,
 - wherein the modified performance strategy comprises one or more reduced power consumption targets for the set of computing systems.
8. The system of claim 7, wherein the control system is further configured to:
 - provide instructions to the set of computing systems to perform the one or more computational operations based on the modified performance strategy.
 9. The system of claim 1, wherein the control system is a remote master control system positioned remotely from the set of computing systems.
 10. The system of claim 1, wherein the control system is a mobile computing device.
 11. The system of claim 1, wherein the control system is configured to receive the power option data while monitoring the set of conditions.
 12. The system of claim 1, wherein the control system is further configured to:
 - provide a request to a qualified scheduling entity (QSE) to determine the power option agreement; and
 - receive power option data in response to providing the request to the QSE.
 13. The system of claim 1, wherein the power option data specify: (i) a first minimum power threshold associated with a first time interval in the set of time intervals, and (ii) a second minimum power threshold associated with a second time interval in the set of time intervals,
 - wherein the second time interval is subsequent to the first time interval.
 14. The system of claim 13, wherein the control system is configured to:
 - determine the performance strategy for the set of computing systems such that the performance strategy comprises:
 - a first power consumption target for the set of computing systems for the first time interval, wherein the first power consumption target is equal to or greater than the first minimum power threshold; and
 - a second power consumption target for the set of computing systems for the second time interval, wherein the second power consumption target is equal to or greater than the second minimum power threshold.
 15. The system of claim 1, wherein a total duration of the set of time intervals corresponds to a twenty-four hour period.
 16. The system of claim 1, wherein the set of conditions monitored by the control system further comprise:
 - a price of power from the power grid; and
 - a global mining hash rate and a price for a cryptocurrency; and
 - wherein the control system is configured to:
 - determine the performance strategy for the set of computing systems based on a combination of at the portion of the power option data, the price of power from the power grid, the global mining hash rate and the price for the cryptocurrency,
 - wherein the performance strategy specifies for at least a subset of the set of computing systems to perform mining operations for the cryptocurrency when the price of power from the power grid is equal to or less than a revenue obtained by performing the mining operations for the cryptocurrency.

US 10,608,433 B1

61

17. A method comprising:
 monitoring, by a computing system, a set of conditions;
 receiving, at the computing system, power option data
 based, at least in part, on a power option agreement,
 wherein the power option data specify: (i) a set of
 minimum power thresholds, and (ii) a set of time
 intervals, wherein each minimum power threshold in
 the set of minimum power thresholds is associated with
 a time interval in the set of time intervals;
 responsive to receiving the power option data, determin-
 ing a performance strategy for a set of computing
 systems based on a combination of at least a portion of
 the power option data and at least one condition in the
 set of conditions, wherein the performance strategy
 comprises a power consumption target for the set of
 computing systems for each time interval in the set of
 time intervals, wherein each power consumption target
 is equal to or greater than the minimum power thresh-
 old associated with each time interval; and
 providing instructions to the set of computing systems to
 perform one or more computational operations based
 on the performance strategy.
 18. The method of claim 17, wherein determining the
 performance strategy for the set of computing systems
 comprises:
 identifying information about the set of computing sys-
 tems; and
 determining the performance strategy to further comprise
 instructions for at least a subset of the set of computing
 systems to operate at an increased frequency based on
 a combination of at least the portion of the power
 option data and the information about the set of com-
 puting systems.
 19. The method of claim 17, further comprising:
 receiving subsequent power option data based, at least in
 part, on the power option agreement, wherein the
 subsequent power option data specify to decrease one
 or more minimum power thresholds of the set of
 minimum power thresholds;

62

responsive to receiving the subsequent power option data,
 modifying the performance strategy for the set of
 computing systems based on a combination of at least
 the portion of the subsequent power option data and at
 least one condition in the set of conditions, wherein the
 modified performance strategy comprises one or more
 reduced power consumption targets for the set of com-
 puting systems; and
 providing instructions to the set of computing systems to
 perform the one or more computational operations
 based on the modified performance strategy.
 20. A non-transitory computer readable medium having
 stored therein instructions executable by one or more pro-
 cessors to cause a computing system to perform functions
 comprising:
 monitoring a set of conditions;
 receiving power option data based, at least in part, on a
 power option agreement, wherein the power option
 data specify: (i) a set of minimum power thresholds,
 and (ii) a set of time intervals, wherein each minimum
 power threshold in the set of minimum power thresh-
 olds is associated with a time interval in the set of time
 intervals;
 responsive to receiving the power option data, determin-
 ing a performance strategy for a set of computing
 systems based on a combination of at least a portion of
 the power option data and at least one condition in the
 set of conditions, wherein the performance strategy
 comprises a power consumption target for the set of
 computing systems for each time interval in the set of
 time intervals, wherein each power consumption target
 is equal to or greater than the minimum power thresh-
 old associated with each time interval; and
 providing instructions to the set of computing systems to
 perform one or more computational operations based
 on the performance strategy.

* * * * *

Exhibit 18

Redacted

Short Message Report


Conversations: 1	Participants: 3
Total Messages: 15	Date Range: 5/4/2019 - 5/9/2019

Outline of Conversations



+ [REDACTED] • 15 messages between 5/4/2019 - 5/9/2019 • Austin Storms <+ [REDACTED] • Austin Storms [REDACTED] • Michael McNamara <+ [REDACTED]

Messages in chronological order (times are shown in GMT -04:00)

 + [REDACTED]

AS Austin Storms <+[REDACTED]> 5/4/2019, 12:57 AM
Storms

MM Michael McNamara <+[REDACTED]> 5/5/2019, 4:00 PM
Storms, great to meet you at the conference

This is me:

MM Michael McNamara <+[REDACTED]> 5/5/2019, 4:00 PM
<https://www.linkedin.com/in/michael-mcnamara-1055211>

AS Austin Storms <+[REDACTED]> 5/5/2019, 4:04 PM
Same here, Michael. I'm not on LinkedIn, but you've got my personal #.

I'll put some feelers out to some of my PM friends this week about what we talked about Fri night. Tty soon.


MM Michael McNamara <+[REDACTED]> 5/5/2019, 4:06 PM
Thanks - that's great

I also think your boxes may have some benefits vs the ones we are doing with JB driver

Lots of stuff to collaborate on

AS Austin Storms <+[REDACTED]> 5/5/2019, 7:43 PM
Absolutely. I can send you specs on the boxes/PDUs/logic design - what's your email?

MM Michael McNamara <+[REDACTED]> 5/5/2019, 7:45 PM
Michael.mcnamara@lancium.com

AS Austin Storms <+[REDACTED]> 5/5/2019, 7:49 PM


MM Michael McNamara <+[REDACTED]> 5/8/2019, 1:45 PM
Storms, can you send me those box design specs please!

AS Austin Storms <+[REDACTED]> 5/8/2019, 4:31 PM
Yep! I'll put it together when I get home tonight

MM Michael McNamara <+[REDACTED]> 5/8/2019, 4:31 PM
Thank you, sir

AS Austin Storms <+[REDACTED]> 5/9/2019, 10:44 AM
Redoing one of the spec sheets for the newer Whatsminer models then emailing over to you

MM Michael McNamara <+[REDACTED]> 5/9/2019, 10:49 AM
Great - thanks

MM Michael McNamara <+[REDACTED]> 5/9/2019, 11:51 AM
Also, have you ever looked at building a GPU box?

AS Austin Storms <+[REDACTED]> 5/9/2019, 11:52 AM
I haven't - but conceptually it's the same. Less electrical load density and less CFM exhaust requirements.

Exhibit 19

UNITED STATES DISTRICT COURT
FOR THE WESTERN DISTRICT OF TEXAS
WACO DIVISION

LANCIUM LLC,

Plaintiff,

v.

LAYER1 TECHNOLOGIES, INC.,

Defendant.

Civil Action No. 6:20-cv-00739

Jury Trial Demanded

COMPLAINT FOR PATENT INFRINGEMENT

Plaintiff Lancium LLC (“Lancium”), by and through its attorneys, brings this action and makes the following allegations of patent infringement relating to United States Patent No. 10,608,433 (“the ’433 patent”). Defendant Layer1 Technologies, Inc. (“Layer1”) infringes the ’433 patent in violation of the patent laws of the United States of America, 35 U.S.C. § 1 *et seq.*

THE PARTIES

1. Plaintiff Lancium LLC is a Limited Liability Company, with its principal office and place of business at 6006 Thomas Road, Houston, Texas, 77041.

2. Upon information and belief, Defendant Layer1 is a Delaware Corporation whose principal office and place of business is at 221 Kearny Street, San Francisco, CA, 94108. Layer1 also has a place of business in Ward County, Texas, which is approximately 100-150 miles west of Midland, Texas, where Layer1 conducts Bitcoin mining operations.

JURISDICTION AND VENUE

3. This is an action for patent infringement arising under the patent laws of the United States of America, Title 35, United States Code.

4. This Court has subject matter jurisdiction over this action pursuant to 28 U.S.C. §§ 1331, 1338(a).

5. This Court has personal jurisdiction over the Defendant Layer1 because Layer1 has committed acts of patent infringement in this District. Upon information and belief, Layer1, utilizes Lancium's patented systems and methods to adjust power consumption based on power option agreements in connection with Layer1's Bitcoin mining operations in this District. Layer1, therefore, has systematic and continuous contacts with this District, regularly transacts business within this District, and regularly and purposefully avails itself of the benefits of this District. This Court further has personal jurisdiction over Layer1 generally because Layer1 maintains a principal place of business in this District (*e.g.*, Layer1's Bitcoin mining facilities are located in this District) and Layer1 regularly conducts business in this District. Layer1, therefore, has established minimum contacts within this District such that the exercise of jurisdiction over Layer1 would not offend traditional notions of fair play and substantial justice.

6. Venue is proper in this District pursuant to 28 U.S.C. §1400(b) because Layer1 has committed acts of patent infringement complained of herein in this District, and has a regular and established place of business in this District.

BACKGROUND

7. Michael McNamara and Raymond Cline, the founders of Lancium, created the company to capitalize on the growth of both renewable energy and distributed computing. The founders understood that the rise of renewable energy would result in greater variability of

electricity production and also electricity price. The growth of more renewable energy would also lead to increased periods and locations of negative-priced energy.

8. Messrs. McNamara and Cline further realized that this increased variability created an opportunity. In 2017, Lancium began work on an entirely new type of data center that could essentially be “turned off” during economically opportune time periods. This new type of data center could operate during periods of negatively-priced or low-priced power and not operate (*i.e.*, not draw load), or operate in a reduced capacity (*i.e.*, draw a limited amount of load), during times when power prices were higher.

9. These “flexible” data centers are useful for many computing workloads, including Bitcoin mining. Lancium developed and patented power management monitoring and control software (“Controllable Load Resource Technology”) that permits data centers to ramp their power consumption up or down in seconds. This technology allows data centers to qualify as Controllable Load Resource(s) (“CLR”) in the Electric Reliability Council of Texas (“ERCOT”). Upon information and belief, Lancium was the first load-only CLR.

10. A load-only CLR can be thought of as interacting with the grid in an inverse fashion to the way that a power generation station interacts with the grid. For example, when the demand for electricity is low (resulting in a low price for electricity), a power generation station would typically decrease its production of power. But in such a situation, a load-only CLR would increase power consumption by, for example, engaging in computationally intensive activities requiring significant amounts of electricity such as Bitcoin mining. Likewise, during times of high-priced electricity (*e.g.*, periods of high electricity demand), a power generation station would typically increase its production of power. A load-only CLR, however, would ramp down (*i.e.*, stop (or reduce) its electricity consumption) by, for example, ceasing (or

reducing) Bitcoin mining operations. When this occurs, the load-only CLR (*i.e.*, flexible data center) receives the difference in the value of the real time electricity versus the data center's pre-existing power purchase agreement price. Lancium's Controllable Load Resource Technology allows, among other things, flexible data centers to operate as load-only CLRs and to participate directly in the energy and ancillary services market.

11. Lancium's Controllable Load Resource Technology also benefits the general public and renewable energy generators. For example, when the load-only CLR ramps down, the electricity that the load-only CLR would have used is available for use by other loads such as consumers (*e.g.*, home owners, businesses, etc.). And when electricity prices are low (*e.g.*, in times of low demand), Lancium's technology permits the flexible data center's computer systems to ramp up to, for example, mine Bitcoins, thereby providing a market for excess electricity generated by renewable energy generators (*e.g.*, wind and solar powered energy generators) helping to mitigate variability and periods and locations with negatively priced energy.

12. Lancium's Controllable Load Resource Technology attracted the interest of many companies, including, upon information and belief, Layer1 as alleged below.

THE PATENT-IN-SUIT


13. Lancium protected its revolutionary technology by, among other things, obtaining patents. On October 28, 2019, Lancium filed a provisional application no. 62/927,119 and, shortly thereafter, utility application no. 16/702,931, which duly and legally issued on March 31, 2020 as U.S. Patent No. 10,608,433 ("the '433 patent") titled "Methods and Systems for Adjusting Power Consumption Based on a Fixed-Duration Power Option Agreement." A true and correct copy of the '433 patent is attached as Ex. A.

14. The '433 patent is assigned to Lancium, which owns all right, title, and interest in and to the '433 patent, including the right to assert all caused of action arising under the '433 patent and the rights to remedies for infringement of the '433 patent.

LAYER1'S INFRINGEMENT

15. Layer1 owns and operates Bitcoin mining data centers in, upon information and belief, Ward County, Texas. Ex. B; Ex. C. Layer1 describes these Bitcoin mining facilities as “game changer[s] in Bitcoin mining.” Ex. B, at 1. “Mining Bitcoin is about converting electricity into money,” says Layer1’s CEO, Alex Liegl. Ex. D, at 2. West Texas power is “the cheapest power in the world, at scale.” Ex. E, at 2. Layer1’s average production cost per bitcoin is \$1,000.00, equating to a 90% profit margin at the bitcoin price of \$9,100.00. Ex. D, at 2.

16. Layer1’s Bitcoin mining facilities not only mine bitcoins. Layer1, upon information and belief, recently qualified as a load-only CLR. Ex. F. Upon information and belief, by entering into “demand response” contracts and utilizing what Layer1 characterizes as “its proprietary demand-response software,” these data centers can be tapped in real time to meet peak market demand by shutting down mining operations at a minute’s notice and instead allowing their load to flow onto the grid. Ex. D, at 2. “In summertime when air conditioners in Dallas, Houston, and Austin are going full tilt, Texas electricity prices sometimes surge to nosebleed levels Layer1 will be able to make more money by shutting off its mining machines and allowing the power to flow through its substation to the grid,” says Alex Liegl. Ex. E, at 3. “This is what being a virtual power plant looks like . . . [s]oftware command instantly decreases or increases many megawatts of electricity and #bitcoin hashrate to stabilize public power grids.” Ex. G.



Layer1’s control system, including its “demand-response” software, upon information and belief, infringes Lancium’s ’433 patent as alleged below and in the attached exemplary claim chart (Exhibit H).

COUNT I – INFRINGEMENT OF U.S. PATENT NO. 10,608,433

17. Lancium incorporates by reference and re-alleges all of the preceding paragraphs of this Complaint as if fully set forth herein.

18. Layer1 infringes at least claims 1-3, 6-9, and 11-20 of the ’433 patent literally or under the doctrine of equivalents in violation of 35 U.S.C. § 271(a) by manufacturing, using, offering to sell, selling, and/or importing infringing systems and methods for adjusting power consumption utilized in or by at least Layer1’s Bitcoin mining facilities (“Infringing Products”). Exhibit H, attached hereto, is an exemplary claim chart showing how the systems and methods of the Infringing Products meet every limitation of, and therefore infringe, each of the above-identified claims.¹

19. Lancium has suffered and continues to suffer damages as a result of Layer1’s infringement in an amount to be determined at trial, which, by law, cannot be less than a reasonable royalty, but may also include lost profits, together with interest and costs as fixed by this Court under 35 U.S.C. § 284.

¹ To the extent the claim chart relies on exhibits not identified and attached to this Complaint, the exhibits are identified in and attached to the claim chart.

20. In addition, Layer1's past and ongoing infringement has caused, and continues to cause, Lancium substantial and irreparable harm for which there is no adequate remedy at law unless and until Layer1 is enjoined by this Court.

21. Layer1's infringement is willful. On May 22, 2020, Lancium notified Layer1 that Layer1's "data center demand response functionality . . . may infringe one or more Lancium's patents." Ex. I. A copy of the '433 patent was attached to the May 22, 2020 email. *Id.* After receiving no response from Layer1, Lancium reached out again on June 2, 2020. Ex. I. Layer1 did not respond, and has not responded to date.

22. Layer1, therefore, has had actual knowledge of the '433 patent since at least May 22, 2020.

23. Upon information and belief, Layer1 continued using Lancium's patented technology to operate its Bitcoin mining facilities in at least West Texas after May 22, 2020 (and continues those operations today). Therefore, at least as of May 22, 2020, and thereafter, Layer1's infringement was willful.

24. Lancium is entitled to a finding of willfulness and enhanced damages under at least 35 U.S.C. § 284 based upon Layer1's willful infringement of the '433 patent.

PRAYER FOR RELIEF

WHEREFORE, Lancium respectfully requests that this Court find in its favor and against Layer1, and that the Court grant Lancium the following relief:

A. Judgment in favor of Lancium that Layer1 has infringed, either literally and/or under the doctrine of equivalents, one or more claims of the '433 patent;

B. An award of all damages adequate to compensate Lancium for Layer1's infringement of the '433 patent;

C. Judgment that Layer1's infringement was willful and that the Court award treble damages for the period of such willful infringement pursuant to at least 35 U.S.C. § 284;

D. An award of pre-judgment and post-judgment interest at the maximum rate permitted by law;

E. A finding that this is an exceptional case and awarding Lancium its costs, expenses, disbursements, and reasonable attorney's fees related to Layer1's infringement under 35 U.S.C. § 285 and all other applicable statutes, rules, and common law;

F. A permanent injunction preventing Layer1, its officers, agents, servants and employees, and those person in active concert and participation with any of them, from infringement of one or more claims of the '433 patent or, in the alternative, if the Court finds that an injunction is not warranted, Lancium requests an award of post-judgment royalty to compensate for future infringement;

G. That Lancium be granted all other relief, in law or equity, as the Court may deem just and proper.

JURY TRIAL

Pursuant to Rule 38 of the Federal Rules of Civil Procedure, Lancium hereby requests a trial by jury on all issues so triable.

Dated: August 14, 2020

Respectfully submitted,

BARNES & THORNBURG LLP

By: s/Mark C. Nelson

Mark C. Nelson
Texas Bar No. 00794361
Daniel A. Valenzuela
Texas Bar No. 24067918
BARNES & THORNBURG LLP
2121 N. Pearl Street, Suite 700
Dallas, TX 75201
Email: mnelson@btlaw.com
Email: dvalenzuela@btlaw.com
Telephone: 214-258-4200
Fax: 214-258-4199

*Attorneys for Plaintiff
Lancium LLC*

Exhibit 20

Redacted in its Entirety

Exhibit 21

Redacted

ATTORNEYS' EYES ONLY – HIGHLY CONFIDENTIAL –
SUBJECT TO PROTECTIVE ORDER

**IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE**

BEARBOX LLC and AUSTIN STORMS,)	
)	
Plaintiffs,)	
)	
v.)	C.A. No. 21-534-MN
)	
LANCIUM LLC, MICHAEL T. MCNAMARA,)	
and RAYMOND E. CLINE, JR.,)	
)	
Defendant.)	

**DEFENDANTS' SECOND SUPPLEMENTAL RESPONSE TO
PLAINTIFFS' INTERROGATORY NO. 3**

Pursuant to Rules 26 and 33 of the Federal Rules of Civil Procedure and the Local Rules of this Court, Defendants Lancium LLC, Michael T. McNamara, and Raymond E. Cline, Jr., (collectively, "Defendants"), by their undersigned attorneys, hereby provide the following supplemental response to Plaintiffs BearBox, LLC and Austin Storms' (collectively, "Plaintiffs") Interrogatory No. 3 served on June 9, 2021 as follows:

PRELIMINARY STATEMENT

Defendants' supplemental response to Plaintiffs' Interrogatory No. 3 is made to the best of Defendants' present knowledge, information, and belief. Defendants' investigation of the facts are ongoing, and Defendants reserve the right to supplement or amend these responses pursuant to the Federal Rules of Civil Procedure, the local rules, the Court's Default Standard for Discovery, Including Discovery of Electronically Stored Information ("ESI"), the Court's Scheduling Order (D.I. 35) and Amended Scheduling Order (D.I. 35), and any other applicable orders. Defendants' responses are not admissions, concessions, or waivers as to the existence, relevance, materiality, foundation, or admissibility of any documents or information.

ATTORNEYS' EYES ONLY – HIGHLY CONFIDENTIAL –
SUBJECT TO PROTECTIVE ORDER

8. Defendants object to Definition and Instruction K (“Relate to,” “Relating to,” or Related to”) as overly broad, unduly burdensome, and ambiguous, particularly the portions of the definition “referring directly or indirectly to, dealing with, or in any way pertaining to.” Such descriptions are subjective, unhelpful, and expand the scope of discovery beyond what is appropriate in this case.

9. Nothing contained in any response herein shall be deemed to be an admission, concession, or waiver by Defendants as to the relevance, competency, materiality, foundation, or admissibility of any document or information provided in response to Plaintiffs’ Requests.

10. To the extent any Interrogatory calling for “all,” “each,” or “every” piece of information as being overly broad and unduly burdensome. It is impossible, to represent, even after a reasonable and diligent search, that all, each, and every piece of information or document falling within a description can be or has been assembled.

RESPONSE TO INTERROGATORIES

INTERROGATORY NO. 3:

Describe in detail the development of each invention claimed in the ’433 Patent including the conception, reduction to practice, and any other development activities for each claimed invention. A complete response to this Interrogatory should include an identification, on a claim-by-claim and element-by-element basis, of which of the purported inventors named on the ’433 Patent conceived of the claimed inventions, reduced the claimed inventions to practice, or otherwise contributed to the development of the claimed inventions, including the dates of each of these activities and an identification of all documents or other evidence that support Your contentions.

ANSWER:

[REDACTED]

[REDACTED]

[REDACTED]

ATTORNEYS' EYES ONLY – HIGHLY CONFIDENTIAL –
SUBJECT TO PROTECTIVE ORDER

[REDACTED]

Case 1:21-cv-00534-GBW-CJB Document 165-1 Filed 06/29/22 Page 196 of 209 PageID #: 4660

A series of horizontal black bars of varying lengths and positions, resembling a barcode or a stylized text representation. The bars are arranged in a vertical sequence, with some being solid black and others appearing as thin white lines on a black background. The overall effect is a dense, abstract pattern of horizontal strokes.

Case 1:21-cv-00534-GBW-CJB Document 165-1 Filed 06/29/22 Page 197 of 209 PageID #: 4661

[illegible]

Case 1:21-cv-00534-GBW-CJB Document 165-1 Filed 06/29/22 Page 198 of 209 PageID #: 4662

The image consists of a single, uniform black rectangle that fills the entire frame. There are no discernible features, text, or patterns other than the solid black color.

ATTORNEYS' EYES ONLY – HIGHLY CONFIDENTIAL –
SUBJECT TO PROTECTIVE ORDER

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

ATTORNEYS' EYES ONLY – HIGHLY CONFIDENTIAL –
SUBJECT TO PROTECTIVE ORDER

[REDACTED]

[REDACTED]

[REDACTED] [REDACTED]

[REDACTED]

[REDACTED]

ATTORNEYS' EYES ONLY – HIGHLY CONFIDENTIAL –
SUBJECT TO PROTECTIVE ORDER

Dated: December 23, 2021

BARNES & THORNBURG LLP

/s/ Chad S.C. Stover

Chad S.C. Stover (No. 4919)
1000 N. West Street, Suite 1500
Wilmington, Delaware 19801-1050
Telephone: (302) 300-3474
E-mail: chad.stover@btaw.com

Mark C. Nelson (admitted *pro hac vice*)
2121 N. Pearl Street, Suite 700
Dallas, TX 75201
Telephone: (214) 258-4140
E-mail: mark.nelson@btlaw.com

Adam M. Kaufmann (admitted *pro hac vice*)
Dana Amato Sarros (admitted *pro hac vice*)
One North Wacker Drive, Suite 4400
Chicago, IL 60606
Tel: (312) 214-8319
E-mail: adam.kaufmann@btlaw.com
E-mail: dana.sarros@btlaw.com

*Attorneys for Lancium LLC, Michael T.
McNamara, and Raymond E. Cline Jr.*

ATTORNEYS' EYES ONLY – HIGHLY CONFIDENTIAL –
SUBJECT TO PROTECTIVE ORDER

CERTIFICATE OF SERVICE

Please take notice that the undersigned hereby certifies that on December 23, 2021 a copy of *Defendants' Second Supplemental Response to Plaintiffs' Interrogatory No. 3* was served on all counsel of record by electronic mail:

Andrew C. Mayo
Ashby & Geddes
500 Delaware Avenue, 8th Floor
P.O. Box 1150
Wilmington, DE 19801
Email: amayo@ashbygeddes.com

Benjamin T. Horton
John R. Labbe
Raymond R. Ricordati III
Chelsea M. Murray
Marshall, Gerstein & Borun LLP
233 South Wacker Drive
6300 Willis Tower
Chicago, IL 60606-6357
Email: bhorton@marshallip.com
Email: jlabbbe@marshallip.com
Email: rricordati@marshallip.com
Email: cmurray@marshallip.com

/s/ Chad S.C. Stover
Chad S.C. Stover (No. 4919)

Exhibit 22

IN THE UNITED STATES DISTRICT COURT
IN AND FOR THE DISTRICT OF DELAWARE

- - -

BEARBOX LLC and AUSTIN)
STORMS,)

Plaintiffs,)

vs.) Civil Action No.
21-534-MN-CJB

LANCIUM LLC, MICHAEL T.)
MCNAMARA, and RAYMOND E.)
CLINE, JR.,)

Defendants.)

- - -

Wilmington, Delaware
Friday, April 22, 2022
Motion to Strike Hearing
and
Discovery Dispute Hearing

- - -

BEFORE: HONORABLE CHRISTOPHER J. BURKE, Magistrate Judge

- - -

APPEARANCES:

ASHBY & GEDDES, P.A.
BY: ANDREW COLIN MAYO, ESQ.
and
MARSHALL GERSTEIN BORUN LLP
BY: BENJAMIN HORTON, ESQ., and
JOHN R. LABBE, ESQ.
(Chicago, Illinois)
Counsel for Plaintiffs

<p style="text-align: right;">Page 2</p> <p>1 APPEARANCES: (Continued)</p> <p>2</p> <p>3 BARNES & THORNBURG LLP</p> <p>4 BY: CHAD S.C. STOVER, ESQ.</p> <p>5 and</p> <p>6 BARNES & THORNBURG LLP</p> <p>7 ADAM KAUFMANN, ESQ.</p> <p>8 (Chicago, Illinois)</p> <p>9 and</p> <p>10 BARNES & THORNBURG LLP</p> <p>11 MARK C. NELSON, ESQ.</p> <p>12 (Dallas, Texas)</p> <p>13 Counsel for Defendants</p> <p>14</p> <p>15</p> <p>16</p> <p>17</p> <p>18</p> <p>19</p> <p>20</p> <p>21</p> <p>22</p> <p>23</p> <p>24</p>	<p style="text-align: right;">Page 4</p> <p>1 be presenting today on behalf of Lancium.</p> <p>2 THE COURT: Okay. Great. Thank you.</p> <p>3 And I should say, for the record, Counsel,</p> <p>4 that we're here this afternoon on this teleconference for</p> <p>5 argument on two different motions: The first is</p> <p>6 defendants' motion to strike the trade secret</p> <p>7 misappropriation counts in the second amended complaint,</p> <p>8 which I believe are Counts 3 and 4, as well as a</p> <p>9 discovery dispute motion, and that is plaintiffs' motion</p> <p>10 to compel the defendants to produce documents concerning</p> <p>11 investments in certain data center facilities in Texas.</p> <p>12 I think it will make sense to first</p> <p>13 address the motion to strike, and then address the</p> <p>14 discovery dispute motion after that.</p> <p>15 And so let me turn, first, to counsel with</p> <p>16 regard to the motion to strike. And I know it's</p> <p>17 defendants' motion, so I'll turn to their counsel first.</p> <p>18 Who's going to be speaking for defendants as to this</p> <p>19 motion?</p> <p>20 MR. KAUFMANN: Good afternoon, Your Honor.</p> <p>21 This is Adam Kaufmann on behalf of defendants.</p> <p>22 THE COURT: Okay. Mr. Kaufmann, thank</p> <p>23 you.</p> <p>24 And, Mr. Kaufmann, I think, in the main,</p>
<p style="text-align: right;">Page 3</p> <p>1 THE COURT: Good afternoon, everybody.</p> <p>2 Just for the record, we're here this afternoon in the</p> <p>3 matter of BearBox LLC and Austin Storms versus Lancium</p> <p>4 LLC, et al., at civil action number 21-534-MN-CJB here in</p> <p>5 our court.</p> <p>6 And before we go further, let's have</p> <p>7 counsel for each side identify themselves for the record.</p> <p>8 We'll start first with counsel for the plaintiff's side,</p> <p>9 and we'll begin there with Delaware counsel.</p> <p>10 MR. MAYO: Good afternoon, Your Honor.</p> <p>11 This is Andrew Mayo from Ashby & Geddes for plaintiffs,</p> <p>12 BearBox and Mr. Storms. I am joined on the telephone</p> <p>13 this afternoon by my cocounsel from Marshall Gerstein &</p> <p>14 Borun. You have Benjamin Horton and John Labbe on the</p> <p>15 line, and Mr. Horton will be presenting the arguments on</p> <p>16 behalf of the plaintiffs today.</p> <p>17 THE COURT: All right. Thank you.</p> <p>18 And we'll do the same for counsel on</p> <p>19 defendants' side, and, again, we'll begin there with</p> <p>20 Delaware counsel.</p> <p>21 MR. STOVER: Good afternoon, Your Honor.</p> <p>22 This is Chad Stover from Barnes & Thornburg, and with me</p> <p>23 are my partners, Mark Nelson from Dallas and</p> <p>24 Adam Kaufmann from Chicago. And Mr. Kaufmann and I will</p>	<p style="text-align: right;">Page 5</p> <p>1 as I understand it, the argument you're making as to why</p> <p>2 these counts should be stricken -- there's various</p> <p>3 arguments, of course. One of them is that the particular</p> <p>4 counts here weren't the type of claims that the Court</p> <p>5 specifically permitted leave to amend on, and that, as a</p> <p>6 result of that, the other side shouldn't be able to amend</p> <p>7 to add trade secret claims at this time.</p> <p>8 But I think it makes sense to focus on</p> <p>9 what I think was kind of the primary line of argument,</p> <p>10 which really goes to kind of the merits or the substance</p> <p>11 of the issue. In other words, even assuming that the</p> <p>12 plaintiffs had made a motion for leave to amend to add</p> <p>13 the new trade secret claims, there, your argument is,</p> <p>14 Look, we're past the deadline for amendment pleadings.</p> <p>15 In fact, we're well past it, and so the movement would</p> <p>16 have to show good cause.</p> <p>17 And that would turn, as an initial matter,</p> <p>18 on the diligence of the movement, and, therefore, it</p> <p>19 would be on the plaintiffs' side to explain why it is</p> <p>20 that they couldn't have added these claims any sooner.</p> <p>21 And that, in turn, leads to an argument</p> <p>22 from you as to whether or not they basically had the</p> <p>23 information relating to these claims at the time of the</p> <p>24 filing of the original complaint. And in response, the</p>

<p style="text-align: right;">Page 18</p> <p>1 THE COURT: Absolutely.</p> <p>2 MR. KAUFMANN: But as to prejudice, I</p> <p>3 mean, you know, yes, there is discovery that we would</p> <p>4 want to take, and the case has continued to progress.</p> <p>5 We're now into the expert discovery phase, and defendants</p> <p>6 have responsive expert reports due in two weeks from</p> <p>7 today. So there really isn't time to take discovery</p> <p>8 before those reports are due, you know, without</p> <p>9 significantly altering the case schedule. So, you know,</p> <p>10 we are at a point where the lack of discovery is</p> <p>11 prejudicial to our case.</p> <p>12 THE COURT: Maybe if I asked it a</p> <p>13 different way. The way that the plaintiff has framed</p> <p>14 what you have said you would need for discovery is that</p> <p>15 you would need three additional interrogatories, five</p> <p>16 additional RFPs, and two additional hours of deposition</p> <p>17 time. Is that an accurate recounting of the</p> <p>18 conversations that you had with them about what discovery</p> <p>19 would look like?</p> <p>20 MR. KAUFMANN: Yes, Your Honor. I believe</p> <p>21 we asked for two hours of additional deposition</p> <p>22 testimony, a few additional requests for admission, and</p> <p>23 the interrogatories. I believe the plaintiffs set it out</p> <p>24 correctly in their responsive paper.</p>	<p style="text-align: right;">Page 20</p> <p>1 assert in your letter in an attempt to demonstrate that</p> <p>2 there is good cause here, and that these are not trade</p> <p>3 secret claims that you knew of, and that you had, and</p> <p>4 that you either could have brought or actually were</p> <p>5 bringing back at the time of the original complaint. I</p> <p>6 think you've said, Look, these trade secrets we're</p> <p>7 referring to in the proposed second amended complaint are</p> <p>8 about specific arbitration-related methods.</p> <p>9 And I think you've cited to the deposition</p> <p>10 that's attached as Exhibit A to the surreply letter, I</p> <p>11 think it's Mr. Storms' deposition on page 309, as a way</p> <p>12 of articulating the specific methods of arbitration that</p> <p>13 are kind of at issue now in the second amended complaint.</p> <p>14 And so that just implicates the question,</p> <p>15 Well, were these methods not in the original complaint?</p> <p>16 And then that implicates the question of, What were the</p> <p>17 trade secrets that were being referred to in the original</p> <p>18 complaint?</p> <p>19 And there, the defendant would say, We</p> <p>20 don't know. And they would say, Judge, you don't know</p> <p>21 because the plaintiff has never told you.</p> <p>22 And so if it's the plaintiffs' burden to</p> <p>23 show good cause and to demonstrate why these new trade</p> <p>24 secrets really are new, that they weren't ones that were</p>
<p style="text-align: right;">Page 19</p> <p>1 THE COURT: Okay. Anything more,</p> <p>2 Mr. Kaufmann, you wish to say about this issue before I</p> <p>3 turn to your colleagues on the other side?</p> <p>4 MR. KAUFMANN: Well, Your Honor, I would</p> <p>5 just reiterate that, you know, I think two key issues</p> <p>6 with this motion are two points of information that</p> <p>7 plaintiffs have just not provided. One, what was the,</p> <p>8 you know, public disclosure of the original trade secret</p> <p>9 claims; and, two, what is the discovery that they claim</p> <p>10 they have received that put them on notice of this new</p> <p>11 trade secret?</p> <p>12 And without either of those pieces of</p> <p>13 information, I'd submit that plaintiffs can't establish</p> <p>14 good cause or diligence in bringing their new claims.</p> <p>15 THE COURT: Okay. Thank you.</p> <p>16 Let me then turn to plaintiffs' counsel.</p> <p>17 And, here, I think it's Mr. Horton who is going to be</p> <p>18 speaking.</p> <p>19 MR. HORTON: That's correct.</p> <p>20 THE COURT: Okay. Mr. Horton, let me jump</p> <p>21 in with questions for you as well, and then I'll</p> <p>22 certainly let you add anything else you wish to add. And</p> <p>23 we'll start and focus mainly on the good cause issue.</p> <p>24 And I think one of the things that you</p>	<p style="text-align: right;">Page 21</p> <p>1 known and being asserted back in the original complaint,</p> <p>2 they've got to tell you that. They've got to give you</p> <p>3 information to help you see that, and they haven't done</p> <p>4 that.</p> <p>5 I guess my fairer question to you there</p> <p>6 would be, Have you told me that? I mean, what were the</p> <p>7 trade secrets that were being referred to in the original</p> <p>8 complaint, and how do I know that they're different from</p> <p>9 the trade secrets that are now said to be at issue with</p> <p>10 the second amended complaint?</p> <p>11 MR. HORTON: Thanks, Your Honor. I guess</p> <p>12 I'll start by saying that the original trade secret</p> <p>13 counts in the first complaint were based on -- I guess</p> <p>14 you can call it an architecture of a system. And shortly</p> <p>15 after filing the complaint, we became aware of a public</p> <p>16 disclosure of that architecture so we didn't feel like,</p> <p>17 in good faith, we could proceed on that trade secret --</p> <p>18 on the trade secret counts on that basis. And so we</p> <p>19 amended the complaint and dropped the trade secret counts</p> <p>20 from it.</p> <p>21 And I'll add, while we're on that</p> <p>22 timeline, Your Honor, that the interrogatory responses</p> <p>23 that defendants are referring to, those interrogatories</p> <p>24 were actually served after the first amended complaint</p>

<p style="text-align: right;">Page 22</p> <p>1 was filed and the trade secret counts were dropped. So</p> <p>2 the trade secrets were no longer at issue in the case</p> <p>3 when that interrogatory responded to.</p> <p>4 But as far as --</p> <p>5 THE COURT: Mr. Horton, wait. I'm sorry.</p> <p>6 Just to back up, you said an architecture of something.</p> <p>7 What was it again?</p> <p>8 MR. HORTON: Yeah. A system architecture,</p> <p>9 Your Honor. So the way the system would be set up where</p> <p>10 different components would be and how they would be</p> <p>11 interconnected.</p> <p>12 THE COURT: But a system architecture</p> <p>13 relating to what?</p> <p>14 MR. HORTON: Related to cryptocurrency</p> <p>15 mining.</p> <p>16 THE COURT: I mean, I'm just trying to --</p> <p>17 obviously, one thing I'm trying to do here is I'm trying</p> <p>18 to understand in my own mind. Like, a key issue here is</p> <p>19 defendants say that the plaintiffs were basically talking</p> <p>20 about the types of arbitrage methods that are referred to</p> <p>21 in the second amended complaint back in the original</p> <p>22 complaint.</p> <p>23 And the plaintiff is saying, No, we</p> <p>24 weren't. No, no, we were talking about something else.</p>	<p style="text-align: right;">Page 24</p> <p>1 THE COURT: And it didn't have anything to</p> <p>2 do with arbitrage methods at all?</p> <p>3 MR. HORTON: That's correct, Your Honor.</p> <p>4 THE COURT: Okay. Now, to that, the other</p> <p>5 side, they cite your response to that interrogatory, and</p> <p>6 I think you've said that the response came after the</p> <p>7 trade secret counts were dropped. But in the</p> <p>8 interrogatory, they used the same phrase "BearBox</p> <p>9 technology" that you used in the original complaint when</p> <p>10 referencing the trade secret.</p> <p>11 And then they say, Look, look at their</p> <p>12 answer. When we asked them what BearBox technology was,</p> <p>13 part of the answer was the sentence beginning on the</p> <p>14 supplementary answer "Mr. Storms also explained." And I</p> <p>15 think they're saying what's being described there is</p> <p>16 energy value arbitrage.</p> <p>17 How come they're wrong? In other words,</p> <p>18 they're suggesting that BearBox technology means the</p> <p>19 kinds of methods for energy value arbitrage they're</p> <p>20 talking about in the second amended complaint, and</p> <p>21 they're citing particular parts of your supplemental</p> <p>22 answer to demonstrate that.</p> <p>23 Why is their reading of your answer</p> <p>24 incorrect.</p>
<p style="text-align: right;">Page 23</p> <p>1 And so incumbent upon that is to say, Here's what we were</p> <p>2 talking about. And you can see it's not the specific</p> <p>3 arbitrage method that we're talking about now in the</p> <p>4 second amended complaint. But I'm still struggling to</p> <p>5 understand what it was you were talking about in the</p> <p>6 first complaint.</p> <p>7 This architecture, what does it have to do</p> <p>8 with -- what more can you tell me about why the two</p> <p>9 things weren't overlapping?</p> <p>10 MR. HORTON: Yeah. Fair enough, Your</p> <p>11 Honor. I'm not doing a good job of explaining that.</p> <p>12 First of all, it had nothing to do with</p> <p>13 energy value arbitrage methods. What it did have to do</p> <p>14 with was more -- when I say "architecture," I'm talking,</p> <p>15 Your Honor, about how different components are connected</p> <p>16 So, in other words, where servers might be in their</p> <p>17 relationship in terms of how they're connected to an</p> <p>18 electricity grid; where the control center may be with</p> <p>19 respect to those servers in the electricity grid.</p> <p>20 Something more high-level, how you set up various</p> <p>21 physical components of a system.</p> <p>22 THE COURT: So it was about certain</p> <p>23 computer architecture?</p> <p>24 MR. HORTON: Yeah.</p>	<p style="text-align: right;">Page 25</p> <p>1 MR. HORTON: It's incorrect, Your Honor,</p> <p>2 because that answer does not discuss or reveal the</p> <p>3 particular method of energy value arbitrage we're talking</p> <p>4 about here, nor would it, because at the time we</p> <p>5 responded to this interrogatory, first of all, no trade</p> <p>6 secret was in the case at all.</p> <p>7 Second of all, we had no reason to</p> <p>8 believe, at that point, that the defendants had</p> <p>9 misappropriated and were using this particular trade</p> <p>10 secret. And the particular trade secret, Your Honor, the</p> <p>11 particular method of energy value arbitrage is very</p> <p>12 specific. It involves specific variables, specific</p> <p>13 estimates or computations about specific performance</p> <p>14 characteristics of particular machines and how those</p> <p>15 values all work together to inform how a system might</p> <p>16 determine at what price to buy energy, at what price,</p> <p>17 when to sell it. It's very specific, Your Honor, and</p> <p>18 that just wasn't contemplated at the time we responded to</p> <p>19 that interrogatory and that information's not in that</p> <p>20 answer.</p> <p>21 THE COURT: I guess when you were</p> <p>22 responding to the interrogatory, did you understand --</p> <p>23 they're obviously asking you to describe all aspects of</p> <p>24 the BearBox technology. Were you attempting -- and the</p>


<p>Page 82</p> <p>1 free to do so, okay?</p> <p>2 All right. So with all that said,</p> <p>3 Counsel, thanks for your patience today. I appreciate</p> <p>4 it. I wish everybody continued health and safety. And</p> <p>5 if there's nothing further, we'll end our teleconference</p> <p>6 today and go off the record. I wish everyone a great</p> <p>7 weekend.</p> <p>8 (The hearing concluded at 4:57 p.m.)</p> <p>9</p> <p>10</p> <p>11</p> <p>12</p> <p>13</p> <p>14</p> <p>15</p> <p>16</p> <p>17</p> <p>18</p> <p>19</p> <p>20</p> <p>21</p> <p>22</p> <p>23</p> <p>24</p>	
<p>Page 83</p> <p>1 CERTIFICATE</p> <p>2</p> <p>3</p> <p>4 I do hereby certify that the foregoing hearing was</p> <p>5 taken before me, pursuant to notice, at the time and</p> <p>6 place indicated; that the statements of participants were</p> <p>7 correctly recorded in machine shorthand by me and</p> <p>8 thereafter transcribed under my supervision with</p> <p>9 computer-aided transcription; that the transcript is a</p> <p>10 true record of the statements made by the participants;</p> <p>11 and that I am neither of counsel nor kin to any party in</p> <p>12 said action, nor interested in the outcome thereof.</p> <p>13</p> <p>14 WITNESS my hand and official seal this 27th day of</p> <p>15 April A.D. 2022.</p> <p>16 </p> <p>17 _____</p> <p>18 Notary Public</p> <p>19</p> <p>20</p> <p>21</p> <p>22</p> <p>23</p> <p>24</p>	

Exhibit 23

Redacted in its Entirety